

O

AR-008-337

DSTO-RR-0071

T

Artificial Intelligence Applications
in Aircraft Systems

Simon Goss and Graeme Murray

S

APPROVED FOR PUBLIC RELEASE

© Commonwealth of Australia

D

Artificial Intelligence Applications in Aircraft Systems

Simon Goss and Graeme Murray

**Air Operations Division
Aeronautical and Maritime Research Laboratory**

DSTO-RR-0071

ABSTRACT

Air Operations Division at the DSTO Aeronautical and Maritime Research Laboratory is developing a capability in the use of Artificial Intelligence (AI), including knowledge based systems technology, in applications related to the operation and support of aircraft systems. A survey of the work program of Air Operations Division was undertaken to identify opportunities offered by advanced computing techniques for the solution of existing research problems. This document describes the findings of the survey. Some of the research opportunities identified have been pursued, and a brief description of progress is provided.

RELEASE LIMITATION

Approved for public release

DEPARTMENT OF DEFENCE

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

19960806 038

Published by

*DSTO Aeronautical and Maritime Research Laboratory
PO Box 4331
Melbourne Victoria 3001*

*Telephone: (03) 9626 8111
Fax: (03) 9626 8999
© Commonwealth of Australia 1996
AR No. AR-008-337
February 1996*

APPROVED FOR PUBLIC RELEASE

Artificial Intelligence Applications in Aircraft Systems

Executive Summary

There are demands on the Australian Defence Force (ADF) and the Defence Science and Technology Organisation (DSTO) for greater performance with fewer resources, and these demands can be expected to continue. Artificial Intelligence (AI), a subfield of computer science, offers the promise of solving technical problems which are otherwise insoluble, or soluble at an unacceptably high cost. In many applications, AI offers much more efficient usage of resources than conventional methodologies.

The Aeronautical and Maritime Research Laboratory supports the ADF through the application of research and development to aircraft and aircraft systems. The work reported here commenced in 1989 in the Aircraft Systems Division (ASD), which maintained a technology base and worked towards solving ADF problems in the areas of:

- aircraft integrated control;
- artificial intelligence;
- avionics and airborne computers;
- mathematical modelling;
- manned simulation;
- human workload analysis; and
- flight management systems.

ASD was integrated into the new Air Operations Division (AOD) in 1992. Within the Division, a capability is being developed in the use of AI, including knowledge based systems technology, in applications related to the operation and support of aircraft systems. In AI the Division has focused on the following areas:

- surrogate operators for combat and training simulators;
- mission management aids;
- support systems for tactical decision making; and
- post processing of simulator data into symbolic summaries.

This document describes the findings of a survey of the work program of the former ASD; the survey was undertaken to identify opportunities offered by advanced computing techniques for the solution of research problems. In focusing on existing research problems, new technical skills can be developed while achieving significant progress in the established work program. In some instances significant progress has been achieved in pursuing these opportunities. A brief description of the research findings is provided.

Authors

Simon Goss

Air Operations Division

Simon Goss, a Senior Research Scientist in AOD, holds BSc (Hons) (1978) and PhD (1985) degrees in Physical Chemistry from La Trobe University and a Graduate Diploma in Knowledge Based Systems (1991) from RMIT. He has worked as an engineer/programmer in industrial control, as a post doctoral research fellow in chemical physics and as a consultant knowledge engineer. He is Chairman of the Victorian Artificial Intelligence Special Interest Group of the Australian Computer Society, and was Convenor of the 4th Australian Conference on Neural Networks, and Chairman of the 2nd Australian Cognitive Science Conference. His research interests lie in knowledge representation, cognitive modelling, and spatial and temporal pattern recognition in air combat modelling, mission planning and human factors analysis applications. He is a member of the Association for Computing Machinery, IEEE, and the Australian Society for Operations Research.

Graeme Murray

Air Operations Division

Graeme Murray, Head of Operations and Performance Analysis in AOD, has an academic background in Physical Chemistry and Pure Mathematics and holds BSc (Melbourne 1966), MSc (Monash 1970) and MA (La Trobe 1979) degrees. He joined the Operational Research Group at the Aeronautical Research Laboratory in 1969 as an analyst, and has wide experience in military operational research and defence issues. He initiated the Aircraft Systems Division research program in the application of artificial intelligence to defence problems in 1989 and leads AOD's research contribution to the Co-operative Research Centre for Intelligent Decision Systems. His research interests include the development of an air operations model using the Australian Artificial Intelligence Institute's dMARS software. He is a member of the Australian Society for Operations Research and a fellow of the Australian Institute of Company Directors.

CONTENTS

	PAGE
1. INTRODUCTION	1
2. THE ARTIFICIAL INTELLIGENCE TOOLBOX	2
3. A PRELIMINARY ASSESSMENT OF ARTIFICIAL INTELLIGENCE FOR AIRCRAFT SYSTEMS APPLICATIONS	6
4. AN EXPERT SYSTEM FOR ESTIMATION OF AIRCRAFT PERFORMANCE	8
5. A SURVEY OF POTENTIAL AI APPLICATIONS IN AIRCRAFT SYSTEMS	11
6. SPECIFIC APPLICATIONS	14
6.1 Speech Recognition for an ATC Simulator	14
6.2 Air Combat Modelling	16
6.2.1 Tactics Selection	18
6.2.2 The Cognitive Model	19
6.3 Aircraft Mission Planning	21
6.3.1 Trajectory Optimisation for a Tactical Strike Aircraft	21
6.4 Tracking Eye Movement	22
7. CENTRE FOR INTELLIGENT DECISION SYSTEMS	25
8. DISCUSSION	28
9. CONCLUSION	30
10. ACKNOWLEDGEMENTS	31
11. REFERENCES	32
APPENDICES	
Appendix 1 An Expert System for Estimation of Aircraft Performance	34
Appendix 2 Report on the Situation Awareness Workshop at IJCAI-91	38
Appendix 3 Agent-Oriented Architecture for Air Combat Simulation	47
Appendix 4 Bibliography	56
FIGURES	
Figure 1. The Artificial Intelligence Toolbox	3
Figure 2. Axes of Complexity in Expert System Development	9
Figure 3. The Cognitive Architecture	20
TABLES	
Table 1. Data Transformation Applications	11
Table 2. Success by Grid Density for Pattern Matcher	23
Table 3. Success by Resolution for Pattern Matcher	23
DISTRIBUTION	
DOCUMENT CONTROL DATA	

1. Introduction

There are demands on the Australian Defence Organisation (ADO) for greater performance with fewer resources, and these demands can be expected to continue. Artificial Intelligence (AI), a subfield of computer science, offers the promise of further progress in problems which are otherwise insoluble, or soluble at an unacceptably high cost. In many applications, AI offers much more efficient usage of resources than conventional methodologies.

The work reported here commenced in 1989 in the Aircraft Systems Division (ASD), which maintained a technology base and worked towards solving ADF problems in the areas of:

aircraft integrated control;	artificial intelligence;
avionics and airborne computers;	mathematical modelling;
manned simulation;	human workload analysis; and
flight management systems.	

ASD was integrated into the new Air Operations Division (AOD) in 1992. In this report, AOD is used generically as including the activities of ASD. Within AOD, a capability is being developed in the use of AI, including knowledge based systems technology, in applications related to the operation and support of aircraft systems. This report covers research projects undertaken between 1989 and 1992 but, where appropriate, statements have been brought up to date to reflect the situation at publication time.

Some difficult technical problems may only be soluble with AI techniques. For example, modern air warfare requires fast and effective decisions by pilots, but the present exponential growth in option lists adds to the pilot's cognitive load, so that some form of computer assistance is highly desirable. Another area of strategic importance where AI may assist is in logistics planning and management in support of military operations, where battlefield requirements change quickly, necessitating extensive preplanning and rapid replanning. AI is a key developing technology with significant potential for assisting and enhancing the effectiveness of the human decision maker. Applications relevant to aircraft systems ranging from avionics fault finding, through mission simulations, to improvement of pilot situation awareness and optimisation of pilot cognitive workload, are central to the work undertaken in support of the ADF.

In AI the Division is focusing its interests in the following areas:

- surrogate operators for combat and training simulators;
- mission management aids;
- support systems for tactical decision making; and
- post processing of simulator data into symbolic summaries.

This document describes the findings of a survey of the work program of the former ASD; the survey was undertaken to identify opportunities offered by advanced computing techniques for the solution of research problems. In focusing on existing research problems, new technical skills can be developed while achieving significant progress in the established work program. In some instances work has been done in pursuing these opportunities, and a brief description of the research findings is provided.

2. The Artificial Intelligence Toolbox

AI can be said to have begun at the Dartmouth Conference in Hanover, New Hampshire, in 1956 (the Dartmouth Summer Research Project on Artificial Intelligence). It is the field of computer science which seeks to understand and implement computer-based technology that can simulate characteristics of human intelligence. Work on constructing AI systems builds on a growing body of widely accepted AI principles and associated computational techniques. A larger scientific goal (to establish an information processing theory of intelligence) underlies much of AI research. A general theory of intelligence is a long term goal for AI researchers, and hence the field is of interest to other researchers such as cognitive psychologists, who are attempting to understand natural intelligence. Currently, much of the international AI research effort is focused on the engineering goal of producing intelligent machines (Smith 1990).

Introductory texts and references in the field of AI include Rich (1983), Winston (1984), and Charniak and McDermott (1985). The recent two volume work, *Artificial Intelligence at MIT* (Winston and Shellard 1990), gives a representative sampling of AI activity. These volumes divide the subject of artificial intelligence into its major areas of application, namely natural language processing, automatic programming, robotics, etc. Nilson(1982) is organised by the general computational concepts involved in the data structures used, the types of operations performed upon these data structures, and the properties of the control strategies used by AI systems.

In applying AI to research problems in AOD we are generally concerned with the *functionality* of intelligent behaviour rather than faithful representation of the internal mechanisms of natural intelligence.

The classical approach in AI is to cast intelligent activity as problem solving with search and representation as the cornerstones, as indicated in Figure 1. *Search* is concerned with the efficient traversal of search space. *Representation* is concerned with recasting problems into appropriate formalisms to facilitate solution and constrain the size of the search space.

Some of the current fields of study are:

(1) Perception (including both vision and speech)

Early efforts at simulating simple, static, visual perception concentrated on statistical pattern recognition and flexible image-understanding systems, although now only the latter falls into the domain of AI. Perceptual tasks are difficult to simulate since they involve data which are typically noisy and usually contain numerous characteristics which must be perceived quickly. Perceptual performance is typified by short time scales, and is not generally amenable to explanation by introspection.

Progress appears to lie more in the domain of pattern matching rather than in knowledge based systems.

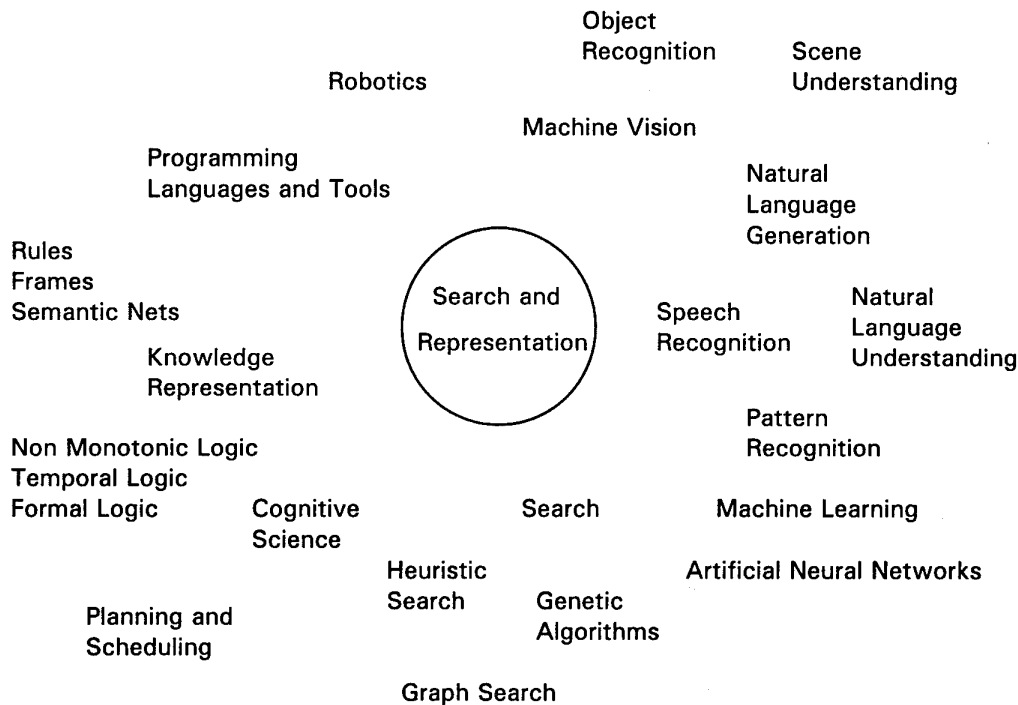


Figure 1 The Artificial Intelligence Toolbox

(2) Natural Language Understanding

Understanding spoken language is a perceptual problem but can be simplified by restricting it to written language. This understanding is still extremely difficult and is usually referred to as natural language understanding (NLU). Understanding a particular topic requires much knowledge about the language, ie its vocabulary and grammar, in addition to the ability to recognise unstated assumptions and implied meanings. The parsers and representations used in NLU are applicable to the codification of time series events and may prove useful in the codification and manipulation of results of studies of human operator performance.

(3) Adaptive Pattern Recognition

The aim of pattern recognition is to build into machines the pattern-information processing capabilities which humans possess. Situations often arise which cannot be described in terms of isolated facts but in terms of *patterns* of interrelated facts. Human perceptive powers are well adapted to pattern processing tasks. For example, we are able to recognise speech utterances and images such as handwriting in a robust manner despite major variations, distortions or omissions. Humans are also able to retrieve information on the basis of associated cues, consisting of only part of a pattern. For example, we are capable of remembering an entire tune after only few notes have been played, or inferring the intention of a pilot in an enemy aircraft by observing a partial track on a radar screen.

An adaptive pattern recognisor is also capable of learning from the environment and modifying its (pattern classification) behaviour in order to improve its performance in that environment. The sub-fields of machine learning and artificial neural networks are aimed at the production of systems which can learn in this way.

(4) Knowledge Representation and Acquisition

Knowledge acquisition includes all activities involved in obtaining information from experts. Knowledge representation is concerned with the process of structuring knowledge about problems in ways to make them easier to solve. Thus where a conventional programmer might seek to reduce a problem in mathematical terms and manipulate it with an algorithm, AI researchers are more interested in expressing knowledge in sentences and pictures which may be manipulated by inferences. Knowledge representation and knowledge acquisition are significant issues in human factors research and the development of air combat models.

An overview of the seminal work in knowledge representation is to be found in Brachman and Levesque (1985). McGraw and Harbison-Briggs (1989) is an accessible summary of research and the pragmatics of methodology in knowledge acquisition.

(5) Machine Vision

The aim of machine vision is to find algorithms for the visual process based on the philosophical assumption that the brain carries out some form of computing when making sense of sensory data. At a pragmatic level the techniques of computational vision can be applied to any set of data which could be viewed as a pixel image. An example is the interpretation of a digital terrain map.

(6) Knowledge Based Systems (KBSs)

These consist of the knowledge about a particular problem domain known as the knowledge base and an inference engine which consists of general problem solving knowledge. *Expert systems* (ESs) are a subset of these that solve problems which are otherwise difficult enough to require significant human expertise for their solution (Buchanan). Of all the commercial activities resulting from AI research, ESs have received the most attention. They now represent a mature technology with applications to problems involving:

- interpretation or the inference of situation descriptions from sensor data;
- prediction;
- diagnosis;
- design or the configuration of objects under constraints;
- monitoring or the comparison of observations with expected results;
- instruction or the diagnosis, debugging and repair of student learning behaviour;
- and control or governing of overall system behaviour.

See Waterman (1986) for more details.

(7) Evidential Reasoning

An evidential reasoning system is a type of knowledge based system which contains rules linking observed evidence for particular findings with hypotheses implied by the evidence to a degree specified by certainty factors. One of the certainty factors relates to the belief in the hypothesis in finding the evidence present while the second indicates how discouraging it is to find the evidence absent. The PROSPECTOR system described in Waterman (1986) is an example of an evidential reasoning system.

(8) Artificial Neural Networks (ANNs)

ANNs are the primary information processing structures of interest in the field sometimes known as neuro-computing. This is the technological discipline concerned with parallel, distributed, adaptive information systems which develop information processing capabilities in response to exposure to an information environment. An important feature of ANNs is their ability to *learn*, which arises from their adaptivity, or ability to change their internal representations as a response to training data (Hecht-Nielsen 1990). Other important attributes, robustness and an ability to deal with noise, are consequences of a distributed internal representation and the ability to generalise as a response to training data.

Some hybrid systems use an ANN to classify input data which then drives a diagnostic KBS. They marry the better characteristics of sub-symbolic and symbolic computation to avoid the restrictions of either method used on its own. (See Hayes 1991, for example.)

(9) Machine Learning

The discipline of machine learning is the study of mechanisms through which intelligent systems improve their performance over time. Research on this topic explores learning in many different domains, employs a variety of methods, and aims for a variety of quite different goals. The field is characterised by its concern with computational mechanisms for learning.

Two dimensions of characterisation are the degree to which the learning process is supervised, and whether the learning is a process of compilation of existing knowledge into a more efficient form or the acquisition of the previously unknown. Learning may also be viewed as a method of summarising or generalising a large number of particular instances into a more compact representation. An example of this would be the summarisation of a large amount of simulator data into a reduced symbolic statement.

For further details of the methods of machine learning see Carbonell (1990) and Shavlik and Dietterich (1990).

3. A Preliminary Assessment of Artificial Intelligence for Aircraft Systems Applications

Conventional computer modelling techniques allow the realistic portrayal of any physical component of a system. High speed digital computers allow high fidelity simulations of complex physical systems, providing an efficient and effective means for studying the behaviour of large systems. Thus computer simulation has become a widely used and accepted tool for studying operational feasibility and operational effectiveness of weapons and equipment. However the behaviour of a human decision maker, a vital factor in many operational effectiveness studies, has proved to be an extremely difficult thing to model realistically with standard procedural computer languages.

A brief preliminary examination of the ASD research program in 1989 identified a range of promising applications for AI:

- air-to-air combat;
- air-to-surface combat;
- flight performance estimation;
- aircrew task performance;
- surrogate operators in simulation; and
- mission management.

In air-to-air combat, the skill of pilots is in their knowledge of the aircraft and weapons they control and the discovering, learning and practising of tactical doctrine applicable to the range of situations which they expect to encounter. Simulation of the processes involved requires a suitable representation of the imprecise information available to each pilot and his/her understanding of the situation. Attempts at representing these processes in procedural computer languages such as FORTRAN, PASCAL and ADA result in very large logic trees, making it difficult to track or modify coded decision processes. Analysis is even more difficult in the case of multiple aircraft engagements, the scenario of greatest interest in modern air combat studies. The use of AI techniques in conjunction with graphical displays may enable pilots to participate in computer modelling procedures, and provide a better representation of air combat performance.

Air-to-surface combat may be studied by means of human-in-the-loop simulation. However, a realistic analysis requires multi-aircraft strikes. It seems unlikely that a sufficient number of realistic pilot stations could be afforded for this work in AOD, so there is a need to generate a number of computer simulated pilots to provide additional aircraft for the multi-aircraft strike force. Providing such surrogates is termed *agency* in simulation, a topic attracting considerable international interest in the international Defence community. AI could provide the means to generate sufficiently authentic surrogate pilots for the study results to be credible to the Service customer.

The flight performance estimation task requires the use of a suite of computer programs in order to calculate the radius of action of aircraft for a specified configuration as well as point performance. Automation of this task would appear to be highly desirable. The initial AI work in ASD was directed towards this goal.

Human factors work involves the study of aircrew task performance as a contribution to manned air mission performance. In order to increase mission effectiveness by redesigning tasks and supporting equipment, it is necessary to know how members of aircrew derive information from sensors, and how this affects their situation awareness. AI techniques may have a role to play in this work.

The AUSTOWER simulator, developed in the time of ASD, provides a training facility for RAAF air traffic controllers (ATCs). It simulates the movements of many aircraft following the instructions of the ATC trainees. In order to make the training realistic, each of these aircraft must appear to move in response to ATC instructions. At present, RAAF staff direct these responses and provide voice traffic. It could be more cost effective if surrogate pilots were provided. This approach has been investigated, and some results are described in this report.

Modern combat aircraft have mission computers which assist the pilot with making tactical decisions by fusing data from a number of sources. This information may be incomplete or contradictory. The process may be assisted by AI techniques in order to avoid excessive or impossible workload demands on aircrew. Some aspects of automated mission planning have also been investigated and are described in Section 6.3.

4. An Expert System for Estimation of Aircraft Performance

AMRL's first involvement in AI was in the Interactive Fault Diagnosis and Isolation System (IFDIS) development for TF-30 engines in RAAF F-111C aircraft. IFDIS, a fault diagnosis advisory system, is an expert system with a knowledge base from the TF-30 maintenance documents and experiential knowledge of TF-30 maintenance experts. It is a user friendly PC based maintenance tool which explains its line of reasoning if required. This work was carried out under the scientific and engineering direction of Propulsion Branch (Larkin et al. 1987), now a part of Airframes and Engines Division (AED).

The first application of AI in the former ASD involved the development of an expert system to assist in the operation of a suite of computer programs which estimate aircraft mission performance. It was believed that an expert system could substantially replace the effort of DSTO specialist staff, by capturing the necessary expertise in a set of rules, so that RAAF staff could operate the generalised mission performance model with minimal DSTO assistance.

The ADF defence effectiveness may be significantly enhanced by a knowledge of the performance of RAAF aircraft when compared with that of likely enemy air weapon systems. There is also a need for mission performance modelling to support RAAF strategic and operational planning, to provide flight manual revisions, and to assist with proposed equipment acquisitions.

The AOD computer programs are used to calculate the radius of action of an aircraft with a specific configuration of stores and weapons, as well as point performance (flight envelope and fan plot). Data are obtained from various sources such as flight manuals and wind tunnel tests. The task of performance estimation is made complex by the variability of both quantity and quality of available input data.

In 1989, DSTO entered into a collaborative arrangement with the Swinburne Institute of Technology (now Swinburne University of Technology) to develop a PC-based advisory expert system to assist in the operation of a suite of computer programs being used to estimate aircraft mission performance. The project was designed to achieve transfer of skills in knowledge engineering and the development of expert systems from Swinburne as part of the process. This acquisition of relevant skills via collaborative work with contracted experts has proven to be effective, particularly in the planning stage of a project. The contractors were responsible for the design, implementation and documentation of the expert system. Deliverables included provision of all necessary software and an evaluation of the system. A fuller description of the work, now completed, is contained in Appendix 1.

The conventional wisdom of knowledge based system development indicates that the best cost/benefit returns are to be found in tasks requiring expertise which is in short supply and/or expensive to retain. The most common problem in constructing an expert system is obtaining sufficient time with the domain expert. That is, the scarcity which motivates the enterprise also makes it difficult to undertake the work. Despite such difficulties, a prototype expert system and draft report were delivered to meet the original time schedule.

It was necessary to formulate the stated problem into a more meaningful overview of the performance task, correctly identifying and sequencing the operations to be carried out. The expert system then guides the user through this process. However, the user needs knowledge of and experience in aeronautical engineering. The prototype system was subjected to a customer evaluation and validation by AOD professional staff with an appropriate level of aeronautical engineering knowledge, and the software refined in the light of comments offered.

The knowledge acquisition required for the construction of an expert system to deal with the entire aircraft performance task is highly labour intensive. To automate the process, a further substantial effort in conventional computer programming would be required to interface the expert system to the existing suite of aircraft performance programs. It was judged that, at least in the short term, the frequency of need for aircraft performance estimates is not sufficient to justify the cost and effort required for full automation. However, the present expert system, based on the VP-Expert¹ shell, is believed to be a useful tool in guiding an aeronautical engineer through the process required to develop the aircraft mission performance estimates.

This task was a valuable learning exercise in AI. The experience confirmed the value of collaborative work in gaining expertise in a new technical area, and showed the benefit of developing a project in discrete stages. The task specified was much more difficult than anticipated, although it has strong similarities to successful applications. The complexity in this application lies in the nature of the engineering expertise in assessing aircraft performance. In particular, judgement based on knowledge of a large number of cases is required. Judgement uses internal subjective mechanisms. It was found that two experts would often arrive at similar determinations by apparently different methods. (Two sources of domain expertise were available to the contractors. These were the expert, and a case study by the expert's predecessor.) Judgements can be rationalised post hoc to give explanation, ie a magistrate in sentencing, but this is more suited to the formalism of case based reasoning as opposed to a knowledge base comprising rules and heuristics.

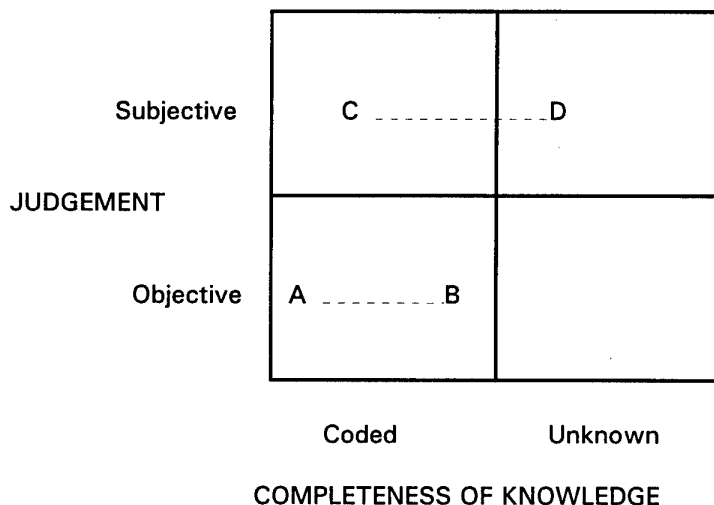


Figure 2. Axes of Complexity in Expert System Development

¹VP-Expert Rule-Based Expert System Development Tool, Paperback Software International, 2830 Ninth Street, Berkeley, California, USA 94710.

Figure 2 illustrates two areas of difficulty in the development of knowledge based systems - judgement and completeness of knowledge. If judgement is objective, ie all experts agree, there are simple decision criteria. Subjective judgement involves weighing evidence and balancing conflicting indications.

If knowledge sources are in codified form such as manuals or tutoring procedures for novices, then development is easier than in the case in which there is no explicit statement of the domain knowledge, requiring the knowledge engineer to create or formalise the knowledge. In attempting to automate the aircraft performance estimation task with an expert system, it was initially believed that the application fell somewhere along the line AB. It became clear in the course of the work that the application lay somewhere along the line CD and was thus a much more difficult task than anticipated.

5. A Survey of Potential AI Applications in Aircraft Systems

In 1990, when the survey was conducted, ASD had six discipline-related groups and 41 tasks. Task types ranged from enabling R&D and skills development, through concept and technology demonstrator development and engineering, to scientific support to the Services.

Two groups specialising in practical avionics and air systems engineering indicated that they saw no use for AI technology either immediately or in the foreseeable future, with the possible exception of avionics diagnostic KBSs. However, the authors believe that KBSs will probably be useful for major mechanical components such as engines and helicopter transmissions. KBSs for avionics systems are most likely to be developed by manufacturers. Deployment of KBSs in avionic instrumentation is unlikely before 1996. Post 1996, KBSs as mission management aids will be potential retrofits to the F/A-18A, and Black Hawk and Seahawk helicopters. An indigenous KBS development capacity is necessary for effective choice and modification of overseas KB equipment to suit local needs. Potential applications in this area have data at the sensor level, ie predominantly numeric rather than symbolic input. Hybrid systems which use a neural net as the front end to provide symbol binding to a KBS have high potential applicability to instrumentation and data fusion applications.

AI technology was seen to be more immediately deployable in support of the activity of ASD research staff involved in activities which were characterised by:

- evaluation and design of the human machine interface in manned systems;
- determining and displaying optimal trajectories to a pilot;
- simulations for training and operational research, and with simulator technologies;
- and commercialisation and systems integration issues.

These activities all tend to involve human performance. AI arose from an attempt to emulate intelligent performance, and some of these activities are consistent with this motivation. The tools available can also be used as data manipulation and transformation techniques. For example, the techniques appropriate to manoeuvre detection in an avionics system, or assisting situation assessment in air combat, may be appropriate to sequence detection in the analysis of simulator data or the analysis of human performance observations. Some other examples are shown in Table 1.

Table 1. Data Transformation Applications of AI

AI Technique	Data Transformation Application
scheduling representation machine learning	sequences codification data summary

The most obvious potential applications of AI to the work program were identified by the preliminary survey reported above. Other potential applications, principally in the facilities and enabling R&D areas, were identified from this survey for further consideration:

- air operations simulation
- flight management & integrated control
- aircrew visual aids research
- aircrew behaviour measurement techniques
- modelling of situation awareness
- human factors in simulation
- analysis of air to air combat
- P-3C avionics systems concepts
- EW displays in tactical aircraft
- AAM target studies
- airborne early warning and control
- aircraft ECM
- ATC tower visual simulator
- F/A-18 data analysis
- F/A-18 SMS software development
- research into operational flight programs
- training simulator technology
- S-70-B2 human factors.

Further applications were identified later, when researchers were beginning to understand the nature of AI work already being undertaken in the Division. Generally the new applications arose out of "wish lists", ie researchers were able to articulate what problems they would like to be able to solve, but were unable to attack through lack of suitable tools. The task plans reported on methods in use to reach an attainable goal rather than some ultimate goal. An example of this is the work on trajectory optimisation for a tactical aircraft. Various non-AI algorithms had been investigated for processing digital terrain data with limited success. Recasting the problem in a state space search formalism allowed the A* (pronounced A-star) search algorithm to be used, and significant progress was made (Section 6.3).

The research program of AOD is focused on practical Service-sponsored tasks, with a small proportion of enabling research supporting a capacity to solve practical problems and offer advice in a timely manner. Real world problems are generally complex and often do not fall squarely within a single computational discipline for their solution. This is well illustrated by the air combat modelling described in a subsequent section of this report. The integration of several computational paradigms is required to achieve a solution, or even a partial solution to a stated problem. Conversely, a single computational approach can offer advantage in several problems which could be said to have elements in common at a higher level of abstraction. For example, state space search is an instance of multi-variate optimisation; the general approach is relevant to generating optimal trajectories from terrain maps, a current investigation, and the exploration of parameter space to determine appropriate manoeuvres to evade a surface-to-air missile. The issues of knowledge representation in modelling tactics selection and situation assessment are also pertinent to the codification and manipulation of information in the cognitive workload studies on the P-3C and Seahawk aircrew.

It became clear when talking with scientific and engineering staff that AI tended to be commonly viewed in one of two ways - either something akin to magic for solving all

problems, or as a technology which had been vastly oversold and had failed to deliver on promises made. It was therefore essential to couple AI investigations of technical problems with an education program for the technical specialists, so that realistic expectations were generated and impossible goals were avoided. Education was achieved by appropriate seminars and continuing technical discussions within ASD. Some of these seminars have been presented by leading AI experts with outstanding international reputations. Successful AI implementations within AOD have contributed to a supportive environment. The stature of AI in AOD has been enhanced by AOD's involvement in the Centre for Intelligent Decision Systems (see Section 7). The invitation to participate arose from an active collaborative research program involving AOD staff.

6. Specific Applications

Several applications were selected for initial development. Progress on these is reported in detail below.

6.1 Speech Recognition for an ATC Simulator

The utility of speech recognition as an operator interface was investigated. The motivation was to examine the utility of speech as a modality of input and output for training simulators in order to reduce the staffing requirements of these simulators. There is also interest in AOD in the potential application of the technologies of speech generation and speech recognition for warning devices and hands-free device switching.

The principal focus of study was the domain of an ATC training simulator. A related application in the domain of ATC was the user interface to an expert system for Flow Control, developed by the Australian Artificial Intelligence Institute (AAIL) for Sydney Kingsford-Smith Airport for the Civil Aviation Authority. Voice operated telephony at Telecom telephone exchanges was investigated by site visits.

The AUSTOWER Air Traffic Control Tower Cabin Simulation portrays traffic moving in real time under a variety of conditions. The simulator includes pilot responses to ATC commands such as "go around", and voice traffic from surrounding sector controllers. The repertoire includes all standard operational and emergency procedures. The method of integrating speech recognition into the simulation is constrained by the method chosen for the generation of imagery. The simulation graphics have pre-computed segments. These join at event nodes where several possible outcomes are available.

At present the simulated pilot/vehicle operator responses are spoken by instructors in another room. The instructors have screens showing alternative actions and a text prompt to the appropriate radio telephony responses.

The problem of removing the extra instructors from this simulation has four components:

- speech recognition (SR)
- natural language understanding (NLU)
- natural language generation (NLG)
- speech generation (SG).

None of these problems is trivial, and there are the further complexities of requiring real time response to continuous speech from a number of speakers. The four areas are inter-related, and mutually constrain design.

The Telecom application is voice operated telephony, and is characterised by a restricted vocabulary with a total vocabulary of about 50 words broken into four sub-vocabularies as opposed to a general ATC application with a vocabulary of 237 words excluding call signs and place names. The AAIL application is a user interface for data entry to an advisory expert system for air traffic flow control at Sydney Airport. Here

too, the vocabulary is restricted compared to general tower cabin ATC, and has a total vocabulary of about 200 words including placenames and common call signs. These have the subtle difference that voice recognition is an auxiliary information channel to existing data input methods. These are well defined problems, more concerned with SR than NLU.

A survey of the state of the art in SR technology was undertaken, and two commercially available systems were evaluated in the laboratory with a senior air traffic controller. Both systems were claimed to have the ability to recognise speaker-dependent continuous speech in real time.

The Votan 2000 system which is used in the Telecom exchanges handles vocabularies of up to 100 words, and employs template matching with limited vocabulary switching. Our own experiments with the system demonstrated significant shortfalls when measured against ATC requirements. AAI experience mirrored our own. SR performance was seriously degraded for a vocabulary having more than twenty words, and there was substantial latency in tracking. Users are constrained to speak in triples (three words at a time with a markedly unnatural rhythm) to avoid problems of co-articulation. It was therefore not feasible to use this system in either ATC application although it adequately copes with the lesser demands of the exchange telephony application.

The Verbex 7000 employs a superior technology in which pattern matching is constrained by a grammar. Co-articulation is addressed in the training method. It copes well with natural rhythms of speech. Recognition levels of up to 97% can be achieved [Castellano 1993a]. This system was found to be technically feasible for the control of simulation as demonstrated for a scaled down ATC tower cabin system [Castellano 1993b]. These findings are corroborated by the successful use of this same technology in currency trading applications (Applied Financial Services Pty Ltd, Melbourne) and a Canadian ATC simulation (ATS Aerospace Inc, Quebec).

The recognition levels when the Verbex system was fully trained were thought to be potentially acceptable for operational use (1993b). However, the arduousness of the training, which took several hours, was unlikely to be acceptable to ATC operators in an operational environment. Reducing the training time to under 90 minutes by way of a different training strategy increased the frustration levels of the subjects ; the training method required tested the patience of a committed and willing operator and was not felt to be practicable for introduction to the workplace. Extra computer processing using message context may increase recognition levels.

These results raise issues in the design of human computer interfaces, the deployment of systems, and training methodology as well as the narrower criteria of the success rates of a technology in recognising utterances. Part of the problem lies in the expectations held for a speech interface. Speech recognition is commonly perceived as simply an exercise in pattern recognition, when the behaviour actually required, that of a natural language interface, is considerably more complex. Some tasks, such as hands-free data entry, require only speech recognition for their solution and these lie comfortably in the domain of the tractable. Those requiring restricted natural language understanding of a constrained domain are at the limit of commercially available technology. Free ranging discourse as seen in science fiction movies is well beyond the current state of the art.

The amount of training necessary for deploying the Verbex 7000 is not in itself a problem in an ATC tower simulator, as several weeks in the training of an air traffic controller are spent in drill rehearsal of dialogue. Mis-recognition in a trained speech recogniser was found experimentally to be due principally to mismatch between the volume in the utterance and the template generated in the training phase. Multiple templates trained at different speech volume improved this performance; however this approach was not considered practicable for complex applications. The utility of the Verbex 7000 as a speech interface for voice control of a simulation was demonstrated with a restricted capacity PC emulation of a tower simulation with the segment architecture used for screen generation in AUSTOWER.

Continuous speech recognition for a wide vocabulary can only be achieved at present with speaker-dependent systems. This has a significant cost in terms of time required to train the system for each user. The most successful commercial products, such as the Verbex 7000, have limited capabilities in natural language understanding; the achieved recognition rates are unacceptably low for applications such as the ATC training simulator. Significant advances in speech recognition system performance can only be achieved via a greatly increased capability in natural language understanding. This will come from the disciplines of computational linguistics, plan recognition and speech-act theory rather than pattern matching. For any given level of technology, speaker-dependent continuous speech recognition will perform more effectively than speaker-independent recognition. A much more complex problem in natural language understanding, that of accepting commands relating to events in the future, has not been considered in any depth. This problem must be solved in order to provide a robust simulation in the ATC training application considered.

6.2 Air Combat Modelling

AOD has a long term interest in modelling air combat in order to carry out effectiveness studies for the RAAF. Current computer simulations are sophisticated physical models. Their weakest aspect is the lack of depth of tactical knowledge in the system. One model is maintained and developed for the study of close combat, in which the opposing aircraft teams remain within sight of each other (20 nautical miles); another model is used for the study of medium range air combat, in which the engagement is largely beyond visual range.

The programs used provide:

- a detailed representation of each aircraft's aerodynamics and propulsion, including altitude and Mach number effects;
- a model for air intercept radar, including clutter effects;
- models for missile dynamics; and
- models for other relevant sensors.

Each tactic considered is represented by a decision tree which is entered as input data for a particular scenario after discussion with the domain experts, fighter pilots. Very large logic trees make it difficult to track or modify the computer coded decision processes. Multiple aircraft engagements, the scenario of greatest interest in modern air combat, compound the problem.

Validation testing of the air combat models is quite difficult. The outputs of sub-system models may be checked against experimental data or against other recognised models. However, confidence in the overall model may only be established by comparison with results from studies in which aircrew have participated.

At present there is no explicit pilot model, partly because of the complexity and variability of pilot behaviour and partly because the main object of the work is to investigate equipments and scenarios rather than pilot skill. Air combat research in AOD has shown that situation assessment and team co-operation, both of which involve human perception, judgement and communication, are critical elements in combat success.

Situation awareness may be defined as the knowledge, understanding, cognition and anticipation of events, factors and variables affecting the safe, expedient and effective conduct of a mission. Situation awareness embraces a number of underlying technologies; it is somewhere between pattern classification and recognition on the one hand and planning and plan-repair on the other. It involves evidential reasoning and models of non co-operative communication. The IJCAI-91 Workshop on Situation Awareness explored these issues. A report on this workshop is given in Appendix 2.

In 1990, four potential contractors were asked to submit research proposals for the modelling of pilot situation awareness in the context of medium range air to air combat. The principal requirement was for a knowledge based system to be developed and tested with a limited data base and restricted set of rules.

It was expected to provide a proof of concept for later development into a more complete representation of the following phases of air combat:

Detection	(pattern recognition);
Classification;	
Recognition/Identification	(perception);
Inference of Intention	(plan recognition);
Threat Assessment;	
Generate Tactical Options	(planning); and
Evaluate and Select Options	(plan repair).

Two research proposals were received from potential contractors, the Australian Artificial Intelligence Institute (AII) and the Swinburne Institute of Technology. Both proposals were considered to have promise, although two quite different methodologies were to be pursued. The AII proposal was attractive in being more immediately applicable, and therefore accorded a higher priority. However, the Swinburne proposal offered the potential to model the variation in performance between individuals. While this was considered to have a greater technical risk, it potentially offered a high return. The two approaches, both of which were funded, are described in more detail in the following sections.

6.2.1 Tactics Selection

AAII proposed a technical approach for studying situation awareness and tactics selection based on six concepts:

- the means of achieving a mission goal, and the activities relevant to achieving that goal can be constructed as a hierarchical network of goals and activities represented procedurally;
- the data store used for holding the information received and intermediate results calculated should be a blackboard architecture²;
- the processes that manipulate the data stored on the blackboard, and the information stored in the hierarchical network should be independent asynchronous processes that embody an understanding of one tactic or one situation;
- the means of fusing information can be based on the Dempster-Shafer theory of evidence (see Shafer 1976);
- when multiple co-operating aircraft are involved, multi-agent distributed reasoning must be used rather than a single centralised system; and
- real-time reasoning methods can be used to allow the tactical knowledge system to attend to the problems of highest priority without ignoring unfulfilled earlier goals.

Time and funding constraints would not allow a full development of AAII's proposal for situation awareness and tactics selection. Consequently, AAII proposed to initiate the work by focusing on tactic representation, selection, and execution tasks, as there are simple ways of emulating situation assessment.

AAII collaborated with AOD in developing a system for representing, selecting, and executing tactics for multiple aircraft engaged in air to air combat. The design was based on SRI International's Procedural Reasoning System (PRS), a multi-agent blackboard system. AAII was responsible for:

- training AOD personnel in the use of PRS;
- design of a multi-agent architecture using PRS;
- setting up the communication protocol between different PRS agents;
- setting up the local knowledge base for each aircraft agent, and the common knowledge base for all aircraft within a group; and
- identifying the co-ordination required between aircraft agents.

AOD was responsible for:

- providing AAII access to air combat expertise;
- developing tactics for each aircraft agent;
- developing global procedures for aircraft co-ordination; and

² A blackboard system is a KBS in which several independent agents communicate via a common database.

- testing and modifying the procedures.

A basic concept demonstrator, including a graphical display of the physics of air combat, was developed to run on a Sun SPARCstation 2. The package allows the tactics to be inspected while the program is running interactively with the display of the physical representation of combat. A small set of aircraft tactics has been coded. The aircraft is represented by a simple point model coded in C.

The concept demonstrator provides a powerful and efficient means for representing tactics in that:

- an operational analyst does not need to program in third generation language source code;
- tactics are separated from the simulation code;
- it is easier to build, modify and display tactics than by conventional coding methods;
- simulated situations are more easily understood as underlying tactics can be displayed; and
- the software can be embedded within the simulation, taking advantage of existing simulation code.

Papers based on this work have been presented at DSTO SIG AI and TTCP meetings, at IJCAI-91 and at AI'92 (see references 9, 18, 20, 21 and 25 in Appendix 5, and references 17, 18, 25 and 26 in Appendix 6). One of the published papers is reproduced in Appendix 3. This collaborative work with AAIL led to DSTO's involvement in the Co-operative Research Centre for Intelligent Decision Systems (see page 23), where the research continues.

It is proposed to expand the repertoire of possible aircraft tactics, replace the graphics module with an X-Windows implementation of the physics of combat, and replace the aircraft point model with a more realistic representation. PRS, the underlying reasoning system, will be replaced by AAIL's dMARS (distributed Multi-Agent Reasoning System) which is written in C++. This will allow the software to be run on a wide range of workstations using Unix. Compatibility with existing AOD air combat models, which run principally on Silicon Graphics machines, will then be achievable.

6.2.2 The Cognitive Model

Swinburne researchers, Professor Iain Wallace and Dr Kevin Bluff, proposed the derivation of an architecture encompassing all of the relevant cognitive processes of a pilot in combat, ie to build a cognitive model of the pilot. This architecture would then provide a framework for a KBS to represent the situation awareness component of air combat modelling, the immediate goal being a computer implementation which is compatible with the computer code available for simulating the dynamics of air combat. A conservative approach was taken; the proposal was split into the description of an architecture, with a second stage leading to a concept demonstrator if the first were deemed successful.

Concentration on psychologically plausible computational models has the desirable outcome of a structure which gives, as side effects, software engineering efficiencies in the doctrinal knowledge base and fuzzy symbol binding. Individual differences in

human performance can be specifically modelled. This approach seemed to be well outside the mainstream of simulation research when initially commenced. However, a presentation by Cmdr D. McBride at the Future Directions in Simulation Symposium Workshop (Melbourne, June 1992) indicated that DARPA was looking at the SOAR architecture (Laird et al. 1987) as the basis for the provision of agency in simulators. BAIRNE (Wallace et al. 1987) is a similar architecture also derived from computational intelligence in psychology. A suitable architecture was derived at Swinburne, including all relevant aspects of a pilot's cognitive processes. It addresses the sequence of phases of air combat simulation. This initial clarification of cognitive processes ensures that a computer implementation will not be limited by previous decisions made on computational grounds.

The sequence of processes involved in situation awareness has close parallels in other situations involving competition between individuals. There are many examples in industry and commerce. Throughout the work the possibility of applying the results to a broader range of contexts was kept in mind in deciding on forms of representation and corresponding implementation methodology. Preliminary work included stepping back from the specific air combat application and looking at the more general problem of decision making in circumstances of constrained information and risk. The IJCAI-91 Workshop on Situation Awareness identified a range of such situations. Several dimensions for the description and characterisation of situation assessment were identified (see Appendix 2).

Figure 3 shows a schematic representation of the model adopted. A hybrid symbolic and non-symbolic representation is used; symbolic and non-symbolic information passes between all modules. The multiple connections indicate the interconnectedness and flow of both types of information. The knowledge base is a semantic network which contains the current state of the system's world knowledge. The environment is perceived by the cognitive model via a sensory apparatus. Actions are taken in the world (environment) by means of the effectors. Stimuli registered from the environment are analysed by non-symbolic feature detectors in the cognitive model. The results are processed to produce symbolic input for semantic nodes of the knowledge base. Selective attention is produced by using the specific symbolic information requirements of semantic nodes as goals for the sensory networks of the cognitive model.




Figure 3. A Cognitive Architecture

This research is reported in international conference papers (see references 8, 11, 17, 23 and 24 in Appendix 5). The next step, now being taken, is to explore the effectiveness in practice of the cognitive architecture and functional mechanisms derived. This is to be achieved via the development of a feasibility prototype which will represent the model developed, but would not be a comprehensive model of pilot situation awareness. A preliminary assessment suggests that a feasibility prototype can be constructed in a fashion suitable for integration with the AAIL model. This preferred approach is being explored. Expansion to a comprehensive alternative simulation of situation awareness for interfacing with existing air combat simulation models is a second possibility.

6.3 Aircraft Mission Planning

Mission planning at the command level is a resource allocation problem. Ideally, a mission planner must generate detailed and sophisticated tactics and be able to provide rapid plan repairs (revisions) during a mission. The time and resources required for comprehensive tactical aircraft mission planning are extensive. It is desirable for the pilot to become a systems manager to achieve mission objectives, rather than being fully occupied in platform control. Research in the Division has focused on the particular problem of determining and displaying optimal trajectories for a tactical aircraft.

6.3.1 Trajectory Optimisation for a Tactical Strike Aircraft

The scenario of interest is that of a tactical strike aircraft on a low-altitude mission involving terrain following/ terrain avoidance. Proximity to the ground enables the pilot to take advantage of the radar shadows of hills and ridge lines so that the aircraft is less visible to ground based radars and weapons systems operators. The problem is to plan an optimal route, given a terrain map, a start point, a target and any mandatory way points. In addition, a capability for replanning en route in response to changed constraints must be provided.

One approach to the problem of terrain avoidance is based on heuristic search of a state space representation. We have used this representation to provide a PC-based research tool for exploring cost functions in the optimisation of trajectories across terrain. This uses the A* algorithm, which combines the better characteristics of branch and bound with dynamic pruning search algorithms. A study of admissible and inadmissible cost functions has produced a counter intuitive result; inadmissible cost functions reduce the search space at the cost of losing a guaranteed optimal solution, but use more computation resources overall. The visualisation tool has been ported to a Sun SPARCstation 2 environment for faster execution and processing larger amounts of data.

A bi-state representation has been developed. In this, manoeuvres are chosen in state space representation, but evaluated in a physical space representation. This approach circumvents problems where discrete approximation of flight paths would lead to a significant accumulation of error (see reference 28 in Appendix 5). This approach has been extended to improve the physical model and increase the number of options for path generation.

A second stream has been pursued directed at the reduction of a terrain map to a connected graph which facilitates search of feature space as well as state space. (A

map can be viewed as the degenerate case of a connected graph in which all points have the trivial relationship of adjacency.) A preliminary investigation of machine interpretation methods for characterisation of terrain has found machine vision techniques to be useful for topographical interpretation of digital terrain maps (see references 26 and 29 in Appendix 5). These techniques can also provide indices for the characterisation of terrain, to determine the appropriate heuristics in the bi-state formulation.

It is proposed to develop and evaluate cost functions grounded in terrain data that reflect world knowledge, and investigate the relationship between the utility of cost functions and terrain type. Other approaches being considered are genetic algorithms for search, and constraint satisfaction programming methods. Research in this area also provides solutions to problems of multivariate optimisation. Other domains of applicability in the AOD research program are optimising crew station layouts and the determination of appropriate counter manoeuvres in parametric exploration of combat effectiveness studies.

6.4 Tracking Eye Movement

A small contract was awarded to the University of Melbourne to explore the practicalities of a pattern matching approach to determining a subject's point of gaze. In current eye tracking work in human factors studies, eye movement monitors of adequate accuracy are also expensive. In that the subject is generally required to wear a helmet mounted video camera which observes the eye, the equipment may affect the process being studied. The pattern matching approach was based on the assumption that we only want to know the geometry in order to determine which region of a display or control panel the subject is looking at. Determination of the geometry is only a means to indicating a region of a display as the point of gaze. This is an indication of the focus of attention, although there is not always a one to one correspondence.

An ANN net-based pattern matcher was trained on a set of examples, in a manner analogous to the way in which a speaker dependent speech recognition system is trained, to indicate which named region is attended to. The training data were images of the eye of the experimental subject obtained with a camera positioned to one side of the centre of the subject's field of regard.

Only two video scan lines from a single colour gun were needed to locate the subject's iris within the centre of the field of view. It is desirable to reduce the amount of image required for processing in order to guarantee real time response of the pattern based eye tracker. If non real-time response is acceptable for a post processing annotation of an experimental record, then this can be increased to allow consideration of more scan lines. The red component of the RGB camera image was found to contain most information.

The proof of concept demonstration was directed at determining the limits of this technology for tracking eyes for a static head with a target fixation area of 200 mm by 200 mm at a viewing distance of 800 mm. A 10 degree field of regard was easily achieved; at 15 degrees performance fell off due to occlusion of the eye in the image by the eyelashes and the iris falling outside the scan lines. At the left and right of extremes of vision, only one edge of the iris crosses the two scan lines. The fixation

area was positioned directly in front of the subject at a distance of 800 mm, appropriate to the situation of a pilot in a cockpit, or operator at a crew station.

Table 2 shows the success rates of this approach for pattern recognisers trained with ten examples for each position on an n-by-n grid. The degradation in performance with increasing grid size is a function of the difficulty in obtaining ten training examples and two test sets of data for each point grid position with an experimental subject in a single extended sitting.

Table 2. Success by Grid Density for Pattern Matcher

Grid Density	3x3	4x4	5x5
Visual Angle (degrees)	7.13	6.67	3.45
Average Training Passes	1000	3000	6000
Average Training Patterns Recognised (%)	99	96.88	93
Maximum Training Patterns Recognised (%)	99	96.88	96
Average Testing Patterns Recognised (%)	82	96.88	56
Maximum Testing Patterns Recognised (%)	89	100	58

Rather than increasing the resolution by increasing the number of points, resolution was tested by starting with a 3x3 grid at 100 mm spacing and decreasing the spacing. Performance fell off sharply between a visual angle of 2.86 and 2.15 degrees. This is in the range of foveal resolution and indicated that this technique could distinguish between attentions to targets separated by 40 mm. The results are reported in Caelli and Wild (1991) and are summarised in Table 3.

Table 3. Success by Resolution for Pattern Matcher

Grid Density (mm)	100	80	60	40	30	20	10	5
Visual Angle (degrees)	7.13	5.71	4.29	2.86	2.15	1.43	0.72	0.36
Average Training Passes	1000	6000	4500	5500	5000	6000	6000	6000
Average Training Patterns Recognised (%)	99	98	97.5	94.5	99	80	41.5	39.5
Maximum Training Patterns Recognised (%)	99	98	100	99	99	80	43	42
Average Testing Patterns Recognised (%)	82	100	83	83	33	28	16.5	13.5
Maximum Testing Patterns Recognised (%)	89	100	94	88	33	28	22	16

This work was successful as a proof in principle of a pattern matching approach to a classification task. It raises the possibility of interleaving a video record of the head up display (HUD) with video images of the pilot. These images of the pilot's head would facilitate the assessment of point of gaze (attribution of region/instrument) by human observers. Pattern matching approaches show considerable promise in human factors

and other categorisation work. The experimental result obtained here has possible applications for determining user attention in video displays.

Any practical application would require estimation or measurement of head pose as well as eye position. (This study used a rigidly constrained head, and measured eye position only.) This could be achieved with a helmet mounted camera, or the use of face marking. User independence is highly desirable. At present, the technique has only been tested for a single subject. Pattern classification offers the potential for video monitoring of a dynamic display for the purpose of real time analysis or initiating some reactive procedure.

7. Co-operative Research Centre for Intelligent Decision Systems

DSTO became involved in the establishment of CIDS in 1991 as a result of technical interaction between AOD staff and the Melbourne AI community, principally at AAIL. This centre is one of the first Co-operative Research Centres (CRCs) set up with Federal Government financial assistance to encourage collaborative research involving private industry, Government departments and academic institutions. CIDS is developing intelligent decision technologies for improving operational management in industry and its supporting infrastructure. Its aim is to aid the growth of value-adding industries for the economic and social benefit of Australia.

The Centre brings together groups with strengths in generic research, applied research and industry. The members of the Centre in 1996 are:

Adacel Pty Ltd

The Australian Artificial Intelligence Institute Ltd (AAIL)

Commonwealth Scientific and Industrial Research Organisation, as represented by the Division of Information Technology (CSIRO)

Computer Power Software Group

Defence Science and Technology Organisation, as represented by the Aeronautical and Maritime Research Laboratory (DSTO)

Royal Melbourne Institute of Technology (RMIT)

The University of Melbourne

Thomson Radar Australia Corporation.

During 1994-95 the Centre concentrated its research and development activities in three areas:

logistics systems

modeling and simulation

document management.

The Centre's program is directed at industrial applications in:

aerospace

energy

manufacturing

minerals

service sector

telecommunicationst

transport

utilities.

Commercialisation of research is to be achieved by focusing in areas of economic importance. The three activities essential to realising the aims of the Centre are:

1. generic research into the design of intelligent systems;
2. the application of that research to problems of substantial economic or social importance; and
3. the commercialisation of those applications.

The following elements are needed in order to achieve these aims:

- research groups of international standing in the relevant research areas;
- organisations experienced in applying research to large scale industrial problems;
- industrial groups able to commercialise the technologies effectively;
- a strongly focused management approach for co-ordinating these activities; and
- integrated training and technology transfer schemes.

Within the Centre, DSTO's research directions have been provided for in CIDS projects in air combat modelling. The Smart Whole AiR Mission Model (SWARMM) is an advanced multi-aircraft military simulation. It is based on dMARS, an advanced real-time reasoning system developed by the Centre. It is a simulation system that will be capable of simulating the physics of whole air missions and the pilot reasoning involved in such missions.

The system will provide DSTO with the ability to evaluate rapidly and test counter-air tactics for the Royal Australian Air Force. It will provide high fidelity simulation of combat aircraft, ground controlled interceptor and surveillance aircraft, advanced reasoning capabilities for modelling the pilot's reasoning process, and sophisticated visualisation tools to enable a better understanding of whole air missions. It will enable DSTO to create rapidly and modify tactics used for the pilot's reasoning process as well as for rapidly setting up a simulation scenario.

The project includes foundational research into the methods and techniques needed for developing the required technology as well as the construction of the simulation software.

AAII and DSTO are the primary Centre participants in this project.

Benefits of the Centre

The cross fertilisation of ideas and integration of techniques is fundamental to the development of intelligent decision systems and the vision of the Centre. CRC funding facilitates:

- major new programs from consolidation of existing research efforts and recruitment of new personnel
- extension and enhancement of existing research programs
- integration of separate but related projects
- stimulation of transfer of ideas between project teams
- generation of new pathways for commercial exploitation

- creation of new facilities for training graduate students and industry personnel.

Benefits for AOD, DSTO and the ADF

- additional resources directed at AOD research interests in air combat and mission planning
- a facility for tackling multi-disciplinary problems with domain experts
- exposure of Defence problems to a wide range of interested researchers
- technologies developed for Defence purposes can be commercialised
- ready access to commercially developed products which may have Defence applications.

8. Discussion

A survey of the research program of AOD identified several opportunities for the application of AI techniques. Those chosen as most promising were:

- air combat modelling;
- speech recognition for an ATC training simulator;
- tactical aircraft mission planning; and
- tracking human eye movement.

Significant progress has been achieved in air combat modelling, where tactics selection has been successfully separated from the code representing the physics of air combat. This achievement is of considerable interest to air combat analysts, offering a more efficient method for studying the relative effectiveness of a range of different tactics. However a more sophisticated model to represent the physics of combat is needed before this approach can be used by the analysts. This need is being addressed collaboratively with AAI within CIDS. An initiative to develop collaboratively a cognitive model of a pilot, in order to facilitate a more realistic representation of pilot behaviour in computer simulations, is proceeding to software implementation. The work on air combat is directed towards simulation of situation awareness, and developments in this area will have a direct impact on manned simulations in AOD in terms of fidelity in simulation. There is a substantial amount of R&D required; representing situation awareness and tactics selection is not simply the implementation of known procedures.

The feasibility of using a commercially available speech recognition product, the Verbex 7000 series, for providing a surrogate pilot to respond to voice commands was demonstrated in the context of an ATC training simulator. This technology is expected to have applicability to research simulators and to other types of training simulators.

A visualisation tool has been developed for an aircraft trajectory optimisation application based on the A* search algorithm. This tool has been provided to the research staff who have been investigating methods for optimal trajectory generation. Alternative algorithms are being investigated in conjunction with a range of cost functions. The tool allows rapid visual assessments to be made. A higher level method of terrain description, based on a machine vision approach, is being explored as a potentially more efficient approach to path selection.

A preliminary study of a non-intrusive way of estimating the point of gaze of a human operator was completed. Suggestions have been made for improving the precision of the estimates. The experimental result has possible applications for determining user attention in video displays. Pattern matching appears to be particularly useful in formulating approaches to human factors and other categorisation work.

AI expertise within AMRL is limited to a small number of people who need to be knowledgeable across the broad spectrum of AI activity in order to identify relevant applications. In addition, they need to know where to locate appropriate external expertise. This is achieved by attending professional meetings and conferences, visiting specialists in particular fields, and using research contracts to solve particular problems. Collaborative research has proven to be very effective in achieving technology transfer to (and from) AMRL personnel. AMRL's involvement in CIDS has

greatly enhanced AOD's capability in exploring and developing AI solutions in its technical research program. This is achieved by way of formal links with technical specialists in other research institutions and CIDS funding for projects of interest to AOD.

9. Conclusion

This report covers research projects undertaken between 1989 and 1992 but, where appropriate, statements have been brought up to date to reflect the situation at publication time. In the research described, AI techniques have allowed significant advances to be made over the results achieved by conventional approaches. These applications are being further developed, and other potential applications for AI are being explored.

As a result of research being pursued, AOD has a clearer picture of the ways in which AI can enhance the Divisional research program. There is a growing awareness within AOD of the contribution which AI can make to some of the research programs. AI paradigms generally address a single aspect of human performance, while many real problems require more than one aspect of human performance to be accounted for in order to obtain a solution. AOD expectations are firmly based on realistic projections from recent experience. In the longer term, relevant AI technology will flow to research areas in AOD to become part of the standard tool kit of techniques familiar to researchers.

10. Acknowledgments

The authors wish to acknowledge the research contributions made by the people listed below:

Aircraft performance expert system

Professor Iain Wallace and Dr Kevin Bluff at Swinburne University of Technology;
The late Michael Howlett at AOD.

Speech Recognition

Dr Andrew Lucas at the Australian Artificial Intelligence Institute (AAIL) and Peter McLeod at the Civil Aviation Authority;
Pierre Castellano and Mario Selvestrel, both formerly at AOD.

Tactics selection in air combat modelling

Dr Andrew Lucas, Dr Anand Rao and David Morley at AAIL;
Serena Steuart at AOD, and Dr Ken Anderson, Ron Muscat, and Mario Selvestrel, all formerly at AOD.

The cognitive model for a pilot

Professor Iain Wallace and Dr Kevin Bluff at Swinburne University of Technology.

Aircraft trajectory optimisation

Dr Grahame Smith, formerly at AAIL and Professor Terry Caelli, formerly at the University of Melbourne;
Mario Selvestrel, formerly at AOD, and Dr Sabrina Sestito at AOD.

Tracking eye movement

Professor Terry Caelli, formerly at the University of Melbourne, and Tom Wild at the University of Melbourne.

11. References

- Brachman R. J. and Levesque H. J. (Eds) 1990,
Readings in Knowledge Representation, Morgan Kaufmann, San Mateo, CA, USA.
- Buchanan B.E. and Duda R.O. 1983,
Principles of Rule-Based Systems, in *Advances in Computers* (Vol 22)
Yovits M.C. (Ed), pp 164-216, Academic Press, New York, USA.
- Caelli T. and Wild T. 1991,
The Feasibility of Using Video Images to Track Human Eye Movement,
Computer Science Department, University of Melbourne, Melbourne. A report
submitted to Aircraft Systems Division, Aeronautical Research Laboratory,
Melbourne, Australia.
- Castellano P. 1994a (in publication),
A Multiple Template Strategy for User-Dependent Continuous Speech Recognition,
DSTO General Document, Air Operations Division, Melbourne, Australia.
- Castellano P. 1994b (in publication),
Airport: A Speech Driven Air Traffic Control Training Simulator, DSTO General
Document, Air Operations Division, Melbourne, Australia.
- Carbonell J. (Ed) 1990,
Machine Learning, MIT Press, Cambridge, MA, USA.
- Charniak, E. and McDermott D. 1985,
Introduction to Artificial Intelligence, Addison-Wesley, Menlo Park, CA, USA.
- Davis E. 1990,
Representations of Commonsense Knowledge, Morgan Kaufmann, San Mateo, CA,
USA.
- Hayes S. 1991,
A Comparison of an Expert System and a Neural Network for Respiratory Monitoring,
RMIT Technical Report, Department of Computer Science, RMIT, Melbourne,
Australia.
- Hecht-Nielson R. 1990,
Neurocomputing, Addison Wesley, Menlo Park, CA, USA.
- Kanal L. and Kumar V.(Eds) 1988,
Search in Artificial Intelligence, Springer Verlag, New York, USA.
- Lairde J.E. , Newell A. and Rosenbloom P.S. 1987,
SOAR: An Architecture for General Intelligence,
Artificial Intelligence, 33, 1-64.

- Larkin M.D. , Frith D.A. and Finlay A.S. 1987,
IFDIS - An Expert System for Diagnosis of Failures in Jet Aircraft Engines,
ARL - Aeropropulsion Technical Memorandum 439, Aeronautical Research
Laboratory, Melbourne, Australia.
- McGraw K. L. and Harbison-Briggs K. 1989,
Knowledge Acquisition: Principles and Guidelines, Prentice Hall, Sydney,
Australia.
- Nilson N. 1982,
Principles of Artificial Intelligence, Springer Verlag, Germany.
- Rich E. 1983,
Artificial Intelligence, McGraw-Hill, Singapore.
- Shafer G. 1976,
A Mathematical Theory of Evidence
Princeton University Press, Princeton, NJ, USA.
- Shavlik J. W. and Dietterich T. G. (Eds) 1990,
Readings in Machine Learning, Morgan Kaufmann, San Mateo, CA, USA.
- Smith R. 1990,
Collins Dictionary of Artificial Intelligence, Harper-Collins, Glasgow, Scotland.
- Wallace L., Klahr D. and Bluff K. 1987,
A Self Modifying Production System Model of Cognitive Development,
in *Production System Models of Learning and Development*,
D. Klahr, P. Langley and R. Neches (eds), MIT Press, Cambridge, MA, USA.
- Watermann D.A. 1986,
A Guide to Expert Systems, Addison-Wesley, Menlo Park, CA, USA.
- Winston P.H. 1984,
Artificial Intelligence (Second Edition), Addison-Wesley, Menlo Park, CA, USA.
- Winston P.H. with Shellard S.A. 1990,
Artificial Intelligence at MIT: Expanding Frontiers, Vols 1 and 2, MIT Press,
Cambridge, MA, USA.

APPENDIX 1

An Expert System for Estimation of Aircraft Performance

The first application of AI in AOD involved the development of an expert system to assist in the operation of a suite of computer programs being used to estimate aircraft mission performance. It was intended that the expert system would substantially replace the effort of DSTO specialist staff, so that RAAF staff could operate the generalised mission performance model with minimal DSTO assistance.

The RAAF defence capability is greatly enhanced by a knowledge of the performance of their aircraft when compared with likely enemy air weapon systems. There is also a need for a mission performance model to support RAAF strategic and operational planning, to provide flight manual revisions and to assist with proposed equipment acquisitions. Task AIR 86/062, which commenced in the 1988-89 financial year, responded to two Air Force Research Requirements, AFRR 10/83, "Development of an F-111C Aircraft Performance Model", and AFRR 10/84, "Development of an F/A-18 Aircraft Performance Model".

Several computer programs were used in task 86/062. A suite of point performance programs developed by the Operational Research Group at AOD (P1, P2 and P4), based on energy-maneuvrability theory, is used to produce detailed manoeuvre plots in various formats. A mission performance model (consisting of the programs MISSION and BERGMAN) calculates the radii of action for various mission profiles and store configurations. PERFORM is an aircraft mission performance program which has been adapted and customised for use with the F-111C. It was expected that the RAAF would continue to require estimation of aircraft performance beyond the three year life of the task. The proposed expert system would assist the RAAF in using the generalised mission performance model.

AOD entered into a collaborative arrangement with the Swinburne Institute of Technology to develop the expert system. The project was designed to achieve transfer of AI skills from Swinburne. This process of acquiring relevant skills via collaborative work with contracted experts is considered to be efficient, particularly in the planning stage of a project. Two major components were required from the contractors:

1. technology transfer to AOD of expert systems technology and knowledge engineering skills; and
2. design, implementation and documentation of the expert system.

The planned deliverables were:

- design of PC-based expert system
- provision of necessary software
- an evaluation of the system
- documentation and user manual.

The specific motivation behind this particular application of AI was AOD's wish to automate the process of constructing estimates of aircraft performance for the RAAF. The contractors initially selected the creation of the basic performance data file AEROPROP.DAT, as described in reference 1, for automation. It was agreed that the specific objective of the project would be the construction of an expert system capable of generating appropriate values for the parameters produced by the estimation programs for inclusion in AEROPROP.DAT. This would involve identifying and modelling the expertise involved in the iterative process of inserting information in the estimation programs and modifying it in the light of evaluation of the acceptability of the parameter values produced.

Expert systems are usually constructed by way of a knowledge acquisition process in which one or more domain experts perform a task while attempting to explain their decision processes to another person, the knowledge engineer. The knowledge engineer then attempts to express these processes as a set of rules. This process has been successful in a number of domains in which an expert's reasoning can be readily expressed as a set of rules. On first impressions, this particular task was amenable to this approach.

In some domains, knowledge acquisition has to become a collaborative process in which the expert and the knowledge engineer jointly develop a conception of the domain being investigated. This systemisation of a domain is achieved by a constructive interaction between the expert and the knowledge engineer. This approach was found to be appropriate for the task being undertaken because the nature of the total expertise being used did not readily translate to a set of rules. In particular, it was discovered that the engineering expertise being used lies in a complex iterative process of generating and evaluating data for inclusion in successive runs of the MISSION program. The knowledge and interpretative expertise is required for cross-validation of program results when initial data are likely to be incomplete and of varying validity (source dependent).

The change in methodology involved a switch from "bottom up" to "top down" approach to the domain structure. The focus was moved from the function of the estimation programs to the purpose of the top level MISSION program. The task being attempted involves expertise in the form of expectations of aircraft performance based on knowledge of their roles, technology levels and relevant design data. The "top down" approach involves initial evaluation of the available data on aircraft, based on expectations of their roles and likely technology levels. Once this is done, the focus moves to the operation of the estimation programs.

The principal features of the software developed are:

1. It provides a means of communication between a knowledge engineer and a domain expert, and a means by which they may jointly generate and accumulate knowledge in the specific domain.
2. Users can conveniently obtain an overview of the variety of data sources and their limitations, the nature of the relational structure of the task and the range of expertise necessary for its performance.

3. An opportunity is provided for users to sample the type of heuristics involved in deriving and evaluating the estimated values required for inclusion in the MISSION program. The structure of the software enables easy modification and augmentation of the repertoire of heuristics employed as increasing experience of task performance enables individual users to derive their own preferred approaches.
4. The software fulfils an auxiliary role in providing guidance to users on the entire series of steps involved in task performance. It also offers a framework for the construction of a more comprehensive auxiliary incorporating heuristics and interpretive assistance.
5. The areas where knowledge creation is required if AOD wishes to automate fully the total aircraft performance assessment task are revealed and a structure provided for their inclusion.

Expert systems are generally believed to be useful for tasks requiring expertise which is in short supply and/or expensive to retain. The most common problem in trying to construct an expert system is obtaining sufficient time with the domain expert. That is, the scarcity which motivates the enterprise also makes it difficult to undertake the work. In this case, the aeronautical engineer whose expertise was required had many other demands on his time, resulting in out of working hours interviews, delays in carrying out parts of the work, and difficulties in co-ordinating the work of all those involved in the process. At one stage the domain expert was absent on duty for a period of eight weeks. Nevertheless, a draft report and prototype expert system was delivered to meet the original time schedule. Reference 1 is the report submitted.

The expert system was subjected to a customer evaluation and validation by AMRL professional staff with an appropriate level of aeronautical engineering knowledge. The contractors modified the software to take account of comments made by such potential users, and prepared a user guide and technical guide for the system (reference 2).

The project was much more difficult than anticipated and highlighted a significant cost/benefit issue. The knowledge acquisition required to construct an expert system to deal with the entire aircraft performance task is highly labour intensive. In addition, a large amount of conventional computer programming would be required to interface the expert system to the existing suite of aircraft performance programs. It was judged that, at least in the short term, the frequency of need for aircraft performance estimates is not sufficient to justify the cost and effort required for full automation. However, the present expert system is judged to be a useful tool in guiding an aeronautical engineer through the process required to develop the aircraft mission performance estimates using the software suite available.

Acknowledgments

Graeme Murray managed the research contract with Swinburne Institute of Technology, working collaboratively with Professor Iain Wallace and Dr Kevin Bluff. The domain expert for the aircraft performance estimation was the late Michael Howlett.

References

1. Rouquerol V.A.E. 1986,
Program MISSION: A Generalized Mission Performance Program, ARL Systems Technical Memorandum 89, Aeronautical Research Laboratory, Melbourne, Australia.
2. Wallace J.G. and Bluff K. 1990,
Automation of the Estimation of Aircraft Performance. A report submitted to Aircraft Systems Division, Aeronautical Research Laboratory. Swinburne Institute of Technology, Melbourne, Australia.
3. Wallace J.G. and Bluff K. 1991,
The Advisor System for Assistance with the Development of Estimation of Aircraft Performance. User Guide and Technical Guide. A report submitted to Aircraft Systems Division, Aeronautical Research Laboratory. Swinburne Institute of Technology, Melbourne, Australia.

APPENDIX 2

Report on the Situation Awareness Workshop at IJCAI-91

Sydney, 24 August 1991

1.0 Introduction

The Situation Awareness Workshop at the 11th International Joint Artificial Intelligence Conference (IJCAI-91) held in Sydney, Australia, arose from a long term commitment at AOD to model air combat.

Present air combat models at AOD are whole system simulations in which there is detailed representation of the physics and sub-systems; aircraft models, weapon systems, radar and electronic counter-measure systems. The element that these models lack is good representation of human decision making. Other groups around the world involved in developing or simulating command and control systems face the same problem. A key aspect of human decision making is the aircrew member's mental picture of the world, often called situation awareness (SA).

One definition of SA is

the knowledge, understanding, cognition and anticipation of events, factors, and variables affecting the safe, expedient and effective conduct of a mission.

The problem of achieving good situation awareness is not confined to military C³I applications. Good situation awareness is as essential for decision making by fireman at the site of a fire, or bankers in foreign currency trading, as for the command of a naval battle group or the mission management of a single tactical aircraft.

The SA workshop sought contributions from military, academic and commercial sectors addressing generic issues of situation awareness. In the call for papers the specific example of air combat was posed and a number of phases identified:-

- a) detection
- b) classification
- c) recognition/identification
- d) inference of intention
- e) threat assessment
- f) generate tactical options
- g) evaluate and select options
- h) execute options

- i) monitor and evaluate effectiveness
- j) iterate on (d).

The similarity of this sequence to commercial activity was highlighted:

"If we substitute the word opportunity for threat, we have a model for business activity. The fourth step, Inference of Intention, is the activity which distinguishes this activity from simpler diagnosis-action systems; there are other agents in the environment."

The scope of approaches sought for the workshop was also indicated:

"Situation awareness embraces a number of underlying technologies. It is somewhere between pattern classification and recognition on the one hand, and planning and plan repair on the other. It involves models of non-cooperative communication. Situation awareness has been tackled as a situated task-specific activity rather than a generic activity. This meeting seeks to provide an opportunity for workers from a range of disciplines and application domain projects to interact and address generic issues."

The late announcement of the acceptance of workshops by the Conference Workshop Chair resulted in a reduction of time available for the call for papers. The papers included were chosen on the basis of the abstract, though published in full in the proceedings. Nine papers were accepted, most came from within DSTO, though several described work done for former employers. There were two sessions of short formal presentations, one chaired by Dr Goss (DSTO-AMRL) who organised and chaired the workshop, the other by Prof Terry Caelli (University of Melbourne) and a summary discussion session chaired by Dr Thomas Strat, a Senior Scientist at SRI International. (Dr Strat is a former US Army officer and holds the rank of Major in the US Army Reserve.)

There were twenty six workshops held in the two days allocated for workshops within IJCAI and a number of ancillary related events such as separate three day conferences or symposia in Logic Programming, Expert Systems, and AI and Design in the days immediately prior to the workshop weekend.

There was some similarity and overlap in the workshop themes, as can be seen from the list of workshop titles in Annex A. Workshop proposers were not told which other workshops had been proposed. In particular the Reasoning in Adversarial Domains Workshop was close in tenor to the SA Workshop. If there had been sufficient notice and communication, the two may have been amalgamated.

Given the number of workshops and related events, the short lead time, the specialised nature of the workshop, the strength of competing events, and the degree of overlap between workshop topics, both the level of attendance (13 registrants) and the quality of papers in the SA workshop was good.

2.0 Analysis

This section attempts to draw an overview of the Workshop, highlight particular points and give an integrated set of conclusions to the broad eclectic mix of papers. Lists of registrants and papers are attached in Annexes B and C respectively.

2.1 Historical Approaches:

There are four possible approaches to the problem of situation awareness:

- a) Ignore it as it is too hard.
- b) Factor it out.
- c) Model the behaviour.
- d) Model the mechanism of the behaviour.

Approach (a), neglect, is the most common in the military science literature. One common variant is to presume the problem has been solved and do planning consequent on having achieved situation awareness of the form of some world description. Another is to do some process of data and information fusion and declare the result of this to be either situation assessment, or the input to the process whereby an operator achieves situation awareness.

Approach (b) has some of the hallmarks of the first approach. The distinguishing element between neglect and discounting is some rationale for neglect other than problem complexity. This is the approach taken by Goss in the Chemical Hazard Response System.

The third approach, behavioural modelling, is most common after neglect. Behavioural modelling has been the impetus for the creation of a number of the formalisms of artificial intelligence. The papers of Gibbon and Aisbett, Rao et al., and Price use or expound methods aimed at reproducing behaviours.

Cognitive modelling, in which the mechanism of cognitive behaviour is modelled as well as the observable behaviour, is a harder row to hoe and not as popular an approach. Bluff et al., and King and Kline address cognitive issues.

2.2 Breakdown by Domain and Method

The papers presented in the workshop covered a range of domains of application:

Surveillance	Gibbon et al., Cosgrove et al., Roney
Air Combat Modelling	Rao et al., Bluff et al.
Robotics	Travers et al.
Chemical Hazard Response	Goss
Human Decision Making	King et al.
Foreign Currency Trading	Price

There was a wider range of approaches to methodology:

Non Monotonic Logic	Gibbon et al.
Belief Desires Intentions Architecture	Rao et al.
Evidential Reasoning System	Goss
Gaussian Numeric Analysis	Roney
Ignore the Facts	King et al.
(Fantasy Based Reasoning)	
Neural Networks	Price

Cognitive Model (Hybrid Architecture)	Bluff et al.
Reactive Planning	Travers et al.

2.3 Dimensions of SA

Examination of the papers and consideration of types of activity characterised by decision making under circumstances of risk, constrained information and time such as air combat, foreign currency trading, fire fighting, and landing an aircraft, even crossing a busy street, yielded a set of dimensions which aid analysis and characterisation.

One gross distinction is our motivation as system builders. Are we seeking to automate a process, simulate a whole system, build a decision aid, improve a human machine interface, or build an autonomous agent?

Performance is another criterion with dimensions of optimality and speed. Ironically, it is easier to produce optimal behaviour, than sub-optimal human-like behaviour.

Several issues arise: What is the speed of response required? Is real-time response required of the computer? Is the human behaviour fast or slow? Does the human use declarative knowledge which he/she is liable to be able to expound, and that will be amenable to symbolic representation? Alternately are we looking at fast procedural knowledge which relies on processes of perception rather than logic?

In terms of tractability of the problem, and overall complexity relatively, problems involving slow logical processes where the knowledge is well defined lie at the easier end of the spectrum; fast cognitive procedural processes where the mechanism or knowledge is not well understood or easily represented lie at the harder end. A related but separate issue is whether the evaluation function used in threat assessment is objective or subjective.

Problem complexity can grow exponentially with small movements along these axes. (Actually the axes in some cases may be substantially non linear and be effectively state transitions across thresholds.)

Other closely related concepts are the degree of risk, the affordability of failure, and the repeatability of failure. Game theorists iterate Prisoner's Dilemma as a tit for tat game. There is a difference in degree of risk *on average* where stochastic considerations apply to the one shot chance. Another perspective is the trade off between the cost of risk exposure and potential benefit gained by action (or inactivity).

The manner in which information flows allows further characterisation. In the Workshop papers, the neural network approach to foreign currency trading (Price) is data driven, while the evidential reasoning system for chemical hazard assessment (Goss) back chains from model based goals.

These dimensions are summarised in Table 1 below.

Table 1: Some dimensions of situation assessment.

Dimension	Extremes	
Cost of Failure	Affordable	Catastrophic
Judgement	Objective	Subjective
Risk	Low	High
"Bites at Cherry"	One	Many
Knowledge	Well Defined	Unknown
Mechanism	Logic	Perception
Performance Required	Satisfy	Optimise
Speed	Slow	Fast
Information Flow	Back Chain	Forward Chain

2.4 Overview of papers

Cosgrove and Hickman describe a problem in the interpretation of data from multiple sources. The work of Gibbon and Aisbett is an evaluation of the utility of non-monotonic logics with this application in mind.

Rao, Georgeff and Sonenberg look at an extension of the belief desires and intention architecture of Georgeff's Procedural Reasoning System to accommodate the notions of joint intentions of agents.

Bluff, Goss and Wallace describe a cognitive architecture which is based upon the earlier work of Bluff and Wallace in cognitive architectures for learning. This hybrid architecture uses both symbolic and sub-symbolic representations. The authors draw distinction between *hot* and *cold* cognitive processes.

King and Kline raise an interesting issue based on psychological work reminding us that human decision making is not necessarily rational, and that decisions can be made about the state of the world in the face of clear evidence to the contrary.

The work of Goss in the assessment of chemical hazard is an application where the problem is essentially that of addressing information already available in a data base, and falls into approach (b) in our taxonomy of approaches to situation assessment, where we can neglect, with justification, both the behaviour and mechanisms of situation assessment and go straight from classification to plan retrieval.

Price's description of the modelling of human performance in foreign currency trading details the application claimed as a fielded system in the advertising of the NeuralWorks Professional software shell. It provides some insight into the nature of decision making and risk acceptance in foreign currency trading, and the essentially pragmatic approach to the acceptance of models in the commercial world.

Roney is concerned with systems where world knowledge can be adequately represented numerically and uses a gaussian analysis to differentiate between the component of uncertainty in measurement.

Travers, Kieronska and Venkatesh have developed a reactive planning system for a robotics application. The robot reacts according to how it perceives the world. Such an action response system is another example of justifiable neglect. There are no high order concepts in the construction of the world picture.

3. Discussion Session Conclusions

Two questions were central to the discussion:

- i. *What is situation assessment?*
- ii. *What are the obstacles?*
 - *What are the yard sticks?*
 - *What should we be working on?*

The conclusions were that:

- a) Situation assessment is a loosely defined term which could embrace cognition and all of the field(s) of artificial intelligence.
- b) To a greater rather than lesser degree the domain of application determines what is important and what are the obstacles to further technical advancement.
- c) Situation assessment requires mechanisms for reasoning under uncertainty.
- d) Situation assessment involves temporal and spatial reasoning.
- e) Situation assessment is the mechanism whereby percepts become beliefs.
- f) Meta-reasoning (reasoning about reasoning processing) is necessary.
- g) Fantasy is a way of expressing excessive persistence in accepting belief.

Annex A of APPENDIX 2
The IJCAI-91 Workshops

1. Parallel Processing for Artificial Intelligence
2. Decision Making Throughout the Decision Process
3. Software Engineering for Knowledge Based Systems
4. Modelling for Intelligent Interaction
5. Artificial Intelligence and Business
6. Integrating Artificial Intelligence and Databases
7. Dynamic Scene Understanding
8. Evaluating and Changing Representation In Machine Learning
9. Fully-Implemented Natural Language Understanding Systems
10. Workshop On Evaluation and Chaos in Cognitive Processing
11. Fuzzy Control Workshop
12. Fuzzy Logic in Artificial Intelligence
13. Intelligent and Cooperative Information Systems:
 Bringing Artificial Intelligence and Information Technologies Together.
14. Natural Language Learning
15. Representing Knowledge in Medical Decision Support Systems
16. Explanation Generation for Knowledge Based Systems
17. Artificial Intelligence in Design
18. Artificial Intelligence Approaches to Production Planning:
 Master Scheduling for Sequencing Tools
19. Objects and Artificial Intelligence
20. Computational Approaches to Non-Literal Language, Metaphor,
 Metonymy, Idiom, Speech Acts, Implicature
21. Theoretical and Practical Design of Rational Agents
22. Reasoning in Adversarial Domains
23. Situation Awareness
24. Computer Vision from Cognitive Science to Industrial Automation
25. Advances in Interfacing Production Systems with the Real World
26. Verification and Validation of Artificial Intelligence Systems

Annex B of APPENDIX 2
Workshop Participants

Dr Kevin Bluff
Swinburne Institute of Technology

Dr Michael King, Psychology
Ballarat University College

Prof Terry Caelli
Computer Science and Psychology
University of Melbourne

Dr Anand Rao
Australian Artificial Intelligence Institute

Dr Victor Ciesielski
Dept of Computer Science
Royal Melbourne Institute of Technology

Dr Richard Price
DSTO - Weapons Systems Research
Laboratory

Dr Stephen Cosgrove
DSTO - Electronics Research Laboratory

Dr Bruce Roney
DSTO - Weapons Systems Research
Laboratory

Dr Greg Gibbon
DSTO - Electronics Research Laboratory

Ms Sabrina Sestito
DSTO - Aeronautical Research Laboratory

Mr Roni Gori
DSTO - Weapons Systems Research
Laboratory

Dr Tom Strat
Artificial Intelligence Centre
SRI International

Dr Simon Goss
DSTO - Aeronautical Research Laboratory

Mr Tony Travers
Dept of Computer Science
Curtin University

Annex C of APPENDIX 2
Situation Awareness Workshop Papers

Strategic Situation Assessment from Uncertain and Low Resolution Data

Stephen Cosgrove and John Hickman
DSTO - Electronics Research Laboratory

Commonsense Logics and Belief Revision in Data Maintenance

Greg Gibbon and Janet Aisbett
DSTO - Electronics Research Laboratory

An Architecture for Situation Awareness

Anand Rao and Mike Georgeff, Australian Artificial Intelligence Institute
and Elizabeth Sonenberg, Computer Science, University of Melbourne

Uncertainty in Situation Assessment

Bruce Roney, DSTO - Weapons Systems Research Laboratory

An Autonomous Agent Planning For Survival in a Complex Environment

Tony Travers, Dorota Kieronska and Svetha Venkatesh
Computer Science, Curtin University

Neural Networks for Foreign Exchange Market Prediction

Richard Price, DSTO - Weapons Systems Research Laboratory
formerly of Bank of America, London

A Cognitive Architecture for Situation Awareness

Kevin Bluff, Swinburne Institute of Technology
Simon Goss, DSTO - Aeronautical Research Laboratory
and Iain Wallace, Swinburne Institute of Technology

A Hazard Action Response System For Chemical Fires and Spills

Simon Goss, DSTO - Aeronautical Research Laboratory
formerly of National Scientific Instrument Training Centre

Imagination in Situational Awareness

Michael King, Psychology, Ballarat University College, and
Colin Kline, Engineering, Ballarat University College

APPENDIX 3

Air combat modelling paper presented at the ARL Future Directions
in Simulation Symposium Workshop, Melbourne, June 1992.

Agent-Oriented Architecture for Air Combat Simulation*

Andrew Lucas, Anand Rao and David Morley

Australian Artificial Intelligence Institute

1 Grattan Street, Carlton, Victoria 3053

Tel: (+61)-3-6637922 Fax: (+61)-3-6637937

Email: *lucas@aaii.oz.au*

Mario Selvestrel and Graeme Murray

Defence Science and Technology Organisation

Aeronautical Research Laboratory

506 Lorimer Street, Fishermens Bend, Victoria 3207

Abstract

Air combat modelling using graphical simulation is a powerful means for development and evaluation of tactics. However, large models are particularly expensive and time-consuming to maintain and modify. Multi-aircraft full mission man-in-the-loop simulators will provide an even more complex programming environment.

Real-time procedural reasoning provides a suitable software environment for the development of an air combat simulation model based on the concept of rational agents. This approach allows the analyst to work at a high level, formulating concepts and aims, while keeping the detailed computer programming hidden.

* This research was supported by the Defence Science and Technology Organisation - Aeronautical Research Laboratory, a *Generic Industry Research and Development Grant* from the Department of Industry, Technology and Commerce, Australia and by the Co-operative Research Centre for Intelligent Decision Systems.

1 Introduction

Computer-based air combat modelling is a powerful tool, widely accepted for its usefulness. The extremely high cost of operating aircraft and their weapons has led to a rapid growth in the development and use of computer simulation models as a basis for tactics development, pilot training, and operational evaluation of weapon systems. However, models having the required level of fidelity are very costly to develop, not easily modified, and lack the ability to show the tactics being used. The analyst has to be able to program in the particular programming language used and the tactics routines are not readily understood by pilots and operational planning staff.

Present simulators combine a high degree of fidelity with the ability to simulate multi-aircraft combat. These typically involve tens of thousands of lines of code and are difficult to manage and modify; combat scenarios involving multiple aircraft are especially difficult to model.

The continuing rapid increase in speed of computing and advances in visual system hardware are allowing further significant advances to be made in simulation capability. However, these advances are dependent on further development and expansion of what is already a large body of software. The challenge is to provide improved performance whilst avoiding the existing shortcomings associated with large simulation codes. These are:

- lack of flexibility (in particular, team tactics are difficult to represent);
- the time and cost associated with maintaining and modifying the simulation code;
- the difficulty of making modifications without introducing errors;
- the need to employ analysts who are also experienced programmers, and are familiar with the model and the code;
- the air combat tactics being embedded in the model; and
- the difficulty in visualising the tactics underlying the model as there is no intuitive means of describing them to pilots and operational personnel.

The aims of the work reported here are to assist the analyst in making use of simulators for developing and evaluating new tactics, and to assist pilots and military planners to evaluate tactics prior to their operational introduction.

These aims must be achieved in a way that makes use of the existing body of simulation software. They can be achieved by embedding the proposed new tactics representation in the current simulation environment, so that modifications required in the code representing the physics of the model are minimal.

A simulation model has a number of functional components:

- a model of the aircraft and weapons physics;
- a model of the aircraft sensors;
- a model of the electronic warfare environment;
- the air combat tactics;
- user interfaces for both pilot and simulation control; and
- a visual system to display the simulation.

The approach taken here relates to the representation, display and execution of air combat tactics by the simulation. The key is to separate the knowledge (tactics) from the physical model and visualisation. This allows modification of the knowledge base without being concerned about the remaining code. Hence, the analyst need only work at the high level, formulating concepts and aims. This approach is eminently suited to dealing with extensive repertoires of pre-determined team tactics. The technology used is known as *real-time procedural reasoning*³.

The current move to use of man-in-the-loop training simulations is accelerating the growth of computer simulation. Multi-aircraft full mission simulators will be perhaps the most demanding military application. High cost will restrict the number of pilot stations available, so that there will be a need to provide computer-generated surrogate pilots for both friendly and enemy aircraft.

2 Situation Awareness and Tactics Representation

The tasks performed by a combat pilot can be broadly divided into two areas -- *situation awareness* and *tactics selection*. Determination of the current situation is called situation awareness and the selection of appropriate actions in response to the situation is called tactics selection. Both stages require sophisticated reasoning and are closely linked. Having determined his current situation, a pilot needs to select and execute certain actions; these actions, in part, determine the next situation.

A pilot's reasoning process can be characterised as consisting of *beliefs* about the real world, *goals* that need to be satisfied, mental *plans* or procedures that satisfy certain goals, and committed partial plans or *intentions* that a pilot has adopted in response to external events or internal goals.

In situation awareness, a pilot must infer the beliefs, goals and intentions of other aircraft from their behaviour. In tactics selection, a pilot must react to his beliefs about the current situation or advance his goals towards a desired situation.

While the problems of situation awareness and tactics selection are difficult for a single pilot in combat with a single enemy, the problems become far more complex when a team (of pilots) is in combat with an enemy team. The team as a whole needs to assess the situation by inferring not only the individual beliefs, goals, and intentions of enemy aircraft, but also the mutual beliefs, joint goals, and joint intentions of the entire enemy team. Similarly, tactics selection by a team is more difficult than the selection of tactics by a single pilot, because of the coordination and synchronisation required.

3 Agent-Oriented Architecture

A rational agent can be viewed as a system situated in the real world, continuously receiving perceptual input and responding by taking actions that affect the world. It can be characterized as having beliefs about the real world, goals that need to be satisfied, compiled plans or recipes for achieving certain goals, and committed partial plans or intentions that have been adopted in response to external events or internal goals. Compiled plans play an important role in such characterizations by side-

³ The term real-time is used here in the context of human behaviour and speed of decision making in a dynamic environment such as air combat, ie., between one second and one minute.

stepping the classical planning problem. More recently, the logical and practical design of such rational agents have received increasing attention.

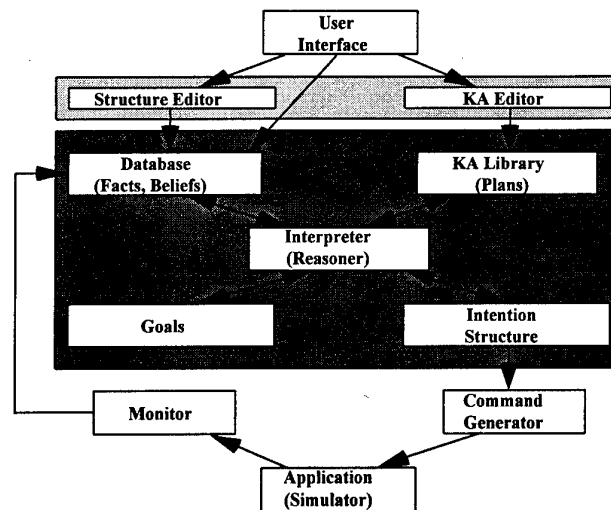


Figure 1: Structure of the Procedural Reasoning System

The Procedural Reasoning System (PRS)⁴ is a situated system based on such an agent-oriented architecture that explicitly reasons about the beliefs, plans, goals, and intentions of a rational agent. The functional modules of the Procedural Reasoning System are given in Figure 1. In PRS, each agent is made up of the following components:

Beliefs: The beliefs of an agent provide information about the state of the environment and are represented as ground first-order logic formulas.

Goals: The goals of an agent are descriptions of desired tasks or behaviours. The goal to achieve a certain formula P is written as $(! P)$; to test for the truth of a formula is written as $(? P)$; and to wait until a formula P is true is written $(^P P)$.

Plan Library: The plan library of an agent consists of a set of plans. Each plan is an abstract specification of possible sequences (and more complex orderings) of actions that can be used to accomplish certain goals or to react to certain situations. Each plan consists of a *name*, an *invocation condition*, which specifies when a plan can be invoked, a *precondition*, which specifies the required belief state of an agent, and a *body*, which describes the possible sequences of actions as a rooted, directed, acyclic graph in which each edge is labelled with a primitive action or a goal expression of the above form. The plan library can also include *metalevel* plans, that is, information about the manipulation of the beliefs, goals, and intentions of the agent itself.

Intention Structure: The intention structure contains a partially ordered set of tasks that the agent is committed to executing. These adopted tasks are called *intentions*. A single intention consists of some initial plan, together with all the sub-plans that are being used in attempting to successfully execute that plan.

⁴ M.P. Georgeff and A.L. Lansky. Procedural knowledge. In *Proceedings of the IEEE Special Issue on Knowledge Representation*, volume 74, pages 1383-1398, 1986.

Intentions may be active or waiting for certain conditions to hold prior to activation. Some of these intentions may be metalevel intentions for decisions on various alternative courses of action.

Interpreter: The interpreter manipulates the above components, selecting appropriate plans based on the system's beliefs and goals, placing them on the intention structure, and executing them. Unless some new belief or goal activates some new plan, the system will try to fulfil any intentions it has previously decided upon. This results in focused, goal-directed reasoning in which plans are expanded in a manner analogous to the execution of subroutines in procedural programming systems. Changes in the environment may lead to changes in the system's goals or beliefs, which in turn may result in the consideration of new plans that are not means to any already intended end. The system is therefore able to change its focus completely and pursue new goals when the situation warrants it.

The system interacts with its environment, including other systems, through its database (which acquires new beliefs and discards old ones in response to changes in the environment) and through the actions that it performs as it carries out its intentions.

4 Team Architecture

Just as a single agent can be considered a rational agent, a team is an organized collection of rational agents. A team can contain other teams as well. This results in a powerful organizational structure capable of modelling hierarchical as well as non-hierarchical social structures. In addition to individual beliefs, goals, plans, and intentions, a team of rational agents will also have: (a) mutual beliefs about the world and about each others' actions; (b) joint goals that need to be achieved collectively; (c) compiled joint plans (common to all team members) that specify the means of satisfying joint goals; and (d) committed joint plans or joint intentions, adopted in response to a joint goal or an external event. Commitment to a joint plan may entail commitment to communicate to other members of the team.

The system for air combat modelling consists of two types of agents -- *team agents* and *aircraft agents*. One aircraft agent is created for each aircraft in the scenario. A team agent is created for each group of aircraft that forms a team. The creation of all these agents, and their interactions with the simulator, are handled by the agent **COORDINATOR**. Team agents and aircraft agents can be dynamically created, destroyed and modified by the **COORDINATOR**.

Beliefs of each aircraft agent include heading and position information about itself and the range and bearing to other aircraft. The beliefs of aircraft agents may vary depending on their role and their situation awareness. For example, an aircraft agent with the role of wingman might have beliefs about the position and heading of the lead aircraft agent. Also, the **COORDINATOR** can selectively pass range and bearing information with respect to other aircraft, thereby simulating the situation awareness of the agents.

A team agent also has beliefs of its own. These beliefs are about the team's heading, position, bearing, etc. Aircraft agents communicate their position and heading information to all their parents.

Aircraft agents tend to have low-level goals and low-level plans. The low-level plans include monitoring the current heading, changing the heading to acquire the desired heading, obtaining a certain lateral separation, flying at a certain heading for a given duration or distance, flying keeping a constant bearing to a given target, firing a missile, guiding a missile, etc. Note that none of the tactics appear as plans for the aircraft agents.

The goals of team agents tend to be higher-level mission goals. As a result, team agents tend to have high-level tactical plans. These plans include the team searching for enemy aircraft in a particular formation, the team doing a pincer or cutoff intercept, the team deciding that a particular sub-team is under threat and doing an evasion, the team choosing a target, etc.

The top-level team adopts goals to perform intercepts, perform evasions, and search for enemy aircraft. Depending on the current scenario the team then chooses an appropriate plan to satisfy one of these goals which then becomes the intention of the team. The successful completion of this intention depends on the lower-level sub-teams adopting certain lower-level intentions. This is achieved by the team sending messages to the lower-level sub-teams to adopt certain intentions. The top-level team intention is then suspended until it receives messages from the lower-level sub-teams that they have accomplished their tasks.

In addition to the goals and intentions adopted as a response to higher-level team intentions the individual aircraft can also adopt goals and intentions of their own to monitor their position, maintain bearing, fire missiles, etc.

5 Sample Procedures

In this section, we illustrate the representational power of the procedures and the ease with which they can be modified.

Evading an enemy aircraft when under threat is quite common in air-to-air combat. The flight path chosen to evade an enemy aircraft depends on a number of factors, some of which are as follows: (a) Has the enemy fired a missile? If so, how long do you have before the missile hits you? (b) Have you fired a missile at the enemy? If so, how long will it take before the missile hits the enemy?

The procedure **Simple Evade**, shown in Figure 2, is invoked when an aircraft has to choose a certain path to evade an enemy aircraft. The context condition specifies that the procedure is to be invoked only when the aircraft is in an evasion mode. If the aircraft does not have a missile flying towards the enemy then it turns 180 degrees from the enemy aircraft. If it does have a missile in the air, then it checks if the enemy also has a missile in the air. If the enemy has a missile in the air, it does a maximum G-turn to radar gimbal limit; otherwise it maintains a constant aspect with the enemy aircraft.

Having tested this simple evasion tactic, let us assume that the analyst wants to modify it so that the evasion tactic can take into account the time left for missile impact. If the time to missile impact is less than 10 seconds the aircraft does a maximum G-turn to radar gimbal limit, if the time is between 10 and 20 seconds the aircraft spirals while maintaining radar lock, and if the time is greater than 20 seconds the aircraft turns 180 degrees from the enemy aircraft. These modifications are made

by adding extra nodes and edges to the procedure **Simple Evade**. The resulting procedure is shown in Figure 3 and is called **Complex Evade**.

Having modified the procedures the analyst can rerun the simulation to obtain different results. Thus, modifying the graphical procedures is equivalent to changing the tactics. Given the modular representation of the procedures in terms of the situations under which they are applicable, the analyst need not worry about the modifications affecting other tactics. This leads to the ease of modification and verification of the tactics represented as graphical procedures. The same approach can be used for representing team tactics, which is investigated in greater detail elsewhere⁵.

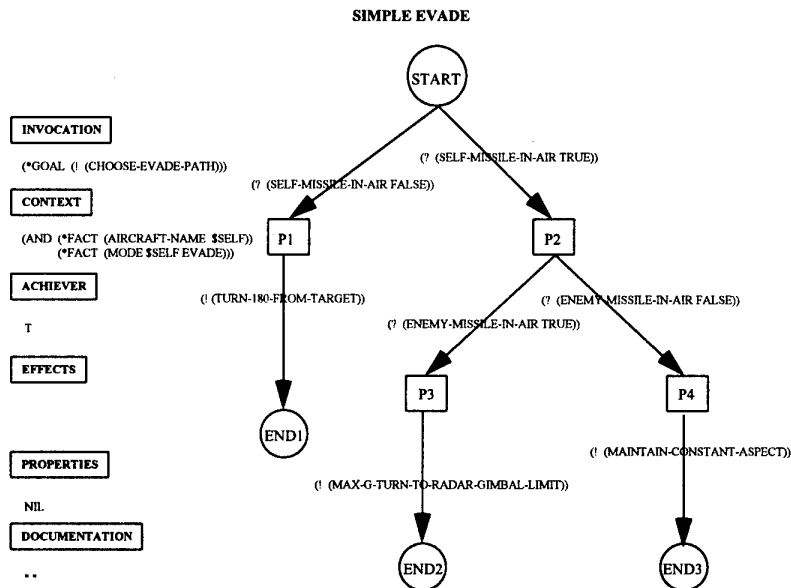


Figure 2: Simple Evasion Procedure

⁵ A. Rao, D. Morley, M. Selvestrel, and G. Murray. Representation, selection, and execution of team tactics in air combat modelling. *Australian Joint Conference on Artificial Intelligence, AI'92*, Hobart, Australia, 1992.

COMPLEX EVADE

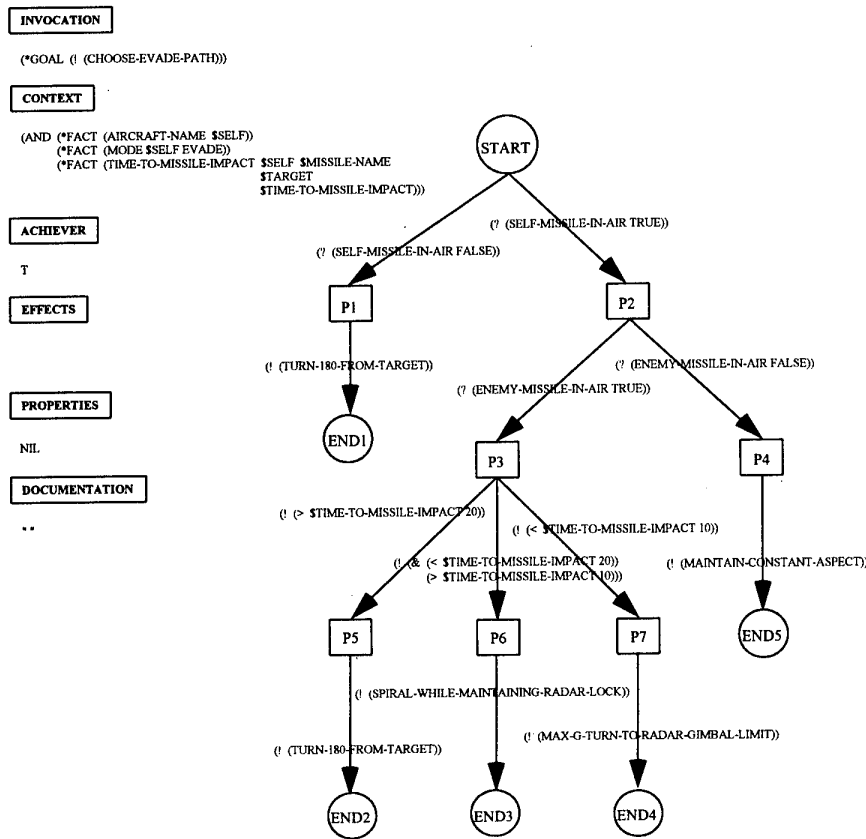


Figure 3: Complex Evasion Procedure

6 Conclusion

The approach outlined here makes it easier (and thus faster) to develop and modify tactics in air combat modelling. The programming complexity of multi-aircraft man-in-the-loop simulators demands such a conceptual breakthrough. Real-time procedural reasoning:

- allows tactics to be constructed and displayed graphically (the analyst does not need to program the tactics in source code);
- separates the tactics from the major body of the simulation code;
- makes it easier to build, modify and display tactics, including team tactics;
- makes simulated situations more easily understood by display of the underlying tactics involved; and
- can be embedded within the simulator, thus taking advantage of existing simulator code.

A prototype of the team tactics representation and execution system, implemented using the Procedural Reasoning System, is running on Sun SPARCstations. PRS is written in Common LISP and an extension of this system, called Multi-Agent Real-time System (MARS), is currently being written in C++. It

will enable direct integration with conventional simulation software. The prototype team tactics system will then be ported to MARS, making it accessible to programmers who are not familiar with LISP.

7 Acknowledgments

The authors acknowledge the contributions of Dr. K. Anderson, Ms. S. Steuart and Mr. R. Muscat during the development of this system.

APPENDIX 4

Bibliography

1. *Using Multi-layered Neural Networks for Learning Symbolic Knowledge*. Conference Paper. S. Sestito and T. Dillon. AI-90, Perth, November 1990. World Scientific, Singapore.
2. *Machine Learning using Single-layered and Multi-layered Neural Networks*. S. Sestito and T. Dillon. IEEE Conference on Tools for AI (TAI-90), Washington DC, USA, November 1990.
3. *Short Term Load Forecasting using an Adaptive Neural Network*, S. Sestito, T. Dillon and S. Leung. *Electrical Power and Energy Systems Journal*. April 1991.
4. *The Use of Sub-symbolic Methods for the Automation of Knowledge Acquisition for Expert Systems*. S. Sestito and T. Dillon. Avignon'91 Conference (11th International Conference on Artificial Intelligence, Expert Systems and Natural Language), Avignon, France, May 1991.
5. *Proceedings of the IJCAI-91 Workshop on Situation Awareness*. (12th International Joint Conference on Artificial Intelligence). S. Goss Editor. Sydney, Australia, August 1991.
6. *A Cognitive Architecture for Situation Awareness*. K. Bluff, S. Goss and J.G. Wallace. *In Proceedings of the IJCAI-91 Workshop on Situation Awareness*. Sydney, Australia. pp. 85-96.
7. *Representation of Team Tactics in Air-to-Air Combat Modelling*. Anand S. Rao, Michael P. Georgeff, and Elizabeth A. Sonenberg, *in Proceedings of the IJCAI- 91 Workshop on Situation Awareness*. Sydney, Australia, August 1991. pp. 498-504.
8. *Using Single-layered Neural Networks for the Extraction of High Level Knowledge Representations*, Sabrina Sestito and Tharam Dillon. *Applied Intelligence Journal*, 1, August 1991. pp. 157-173.
9. *A Cognitive Framework for Hybrid Systems*. J.G. Wallace and K. Bluff. *Proceedings of IJCNN-91 (International Conference on Neural Networks)*, Singapore, November 1991.
10. *Automation of Knowledge Acquisition using Neural Networks*. Sabrina Sestito. PhD thesis. La Trobe University, Melbourne, December 1991.
11. *An Approach Towards the Verification of Expert Systems Using Numerical Petri Nets*. N.K. Liu. *International Journal of Intelligent Systems*, 6, 1991. pp. 255-276.
12. *Formal Description and Verification of Expert Systems*. N.K. Liu. PhD thesis. La Trobe University, Melbourne, December 1991.

13. *Situation Awareness in Air Combat Modelling*. A report submitted to the Aeronautical Research Laboratory. J.G. Wallace and K. Bluff, Swinburne Institute of Technology, Melbourne, May 1992.
14. *Automated Knowledge Acquisition of Rules with Continuously Valued Attributes*. Sabrina Sestito and Tharam Dillon. Avignon'92 Conference, Avignon, France, May 1992.
15. *Agent-Oriented Architecture for Air Combat Simulation*. Andrew Lucas, Anand Rao, David Morley, Mario Selvestrel and Graeme Murray. Future Directions in Simulation Symposium Workshop, Melbourne, June 1992.
16. *Who Saw What: Getting Symbolic Descriptions of Generated Imagery in Research Simulators*. Simon Goss. Future Directions in Simulation Symposium Workshop, Melbourne, June 1992.
17. *Knowledge Representation and Decision Making: A Hybrid Approach*. J.G. Wallace, S. Goss, and K. Bluff. Proceedings of IEA-AIE 92, Paderborn, Germany, June 1992. Springer-Verlag, Germany. pp. 525-528.
18. *Semantic Transparency and Cognitive Modelling of Complex Performance*. J.G. Wallace and K. Bluff. 5th Irish Conference on Artificial Intelligence and Cognitive Science, University of Limerick, Ireland, September 1992. Springer-Verlag, Germany. pp. 71-83.
19. *Representation, Selection and Execution of Team Tactics in Air Combat Modelling*. Anand Rao, David Morley, Mario Selvestrel and Graeme Murray. AI'92, Hobart, November 1992. World Scientific, Singapore. pp. 185-190.
20. *Generating Aircraft Flight Paths using State Space Search*. Poster Paper. M. Selvestrel, S. Goss and G. Smith. AI'92, Hobart, November 1992.
21. *The Detection of Features in Digital Terrain Maps*. Poster Paper. S. Sestito, T. Caelli, and S. Goss. AI'92, Hobart, November 1992.
22. *Knowledge Acquisition of Conjunctive Rules using Multi-layered Neural Networks*. Sabrina Sestito and Tharam Dillon. International Journal of Intelligent Systems, 8, 1992. pp. 779-806.
23. *Modelling Uncertainty in Expert Systems*. N.K. Liu. Proceedings of Pacific Rim International Conference on Artificial Intelligence, Korea, September 1992. IOS Press.

Artificial Intelligence Applications in Aircraft Systems

Simon Goss and Graeme Murray

AUSTRALIA

DEFENCE ORGANISATION

Defence Science and Technology Organisation

Chief Defence Scientist	}	shared copy
FAS Science Policy		
AS Science Corporate Management		
Counsellor Defence Science, London (Doc Data Sheet only)		
Counsellor Defence Science, Washington (Doc Data Sheet only)		
Scientific Adviser to Thailand MRD (Doc Data Sheet only)		
Scientific Adviser to the DRC (Kuala Lumpur) (Doc Data Sheet only)		
Senior Defence Scientific Adviser/Scientific Adviser Policy and Command (shared copy)		
Navy Scientific Adviser (3 copies Doc Data Sheet and one copy of the distribution list)		
Scientific Adviser - Army (Doc Data Sheet and distribution list only)		
Air Force Scientific Adviser		
Director Trials		

Aeronautical and Maritime Research Laboratory

Director AMRL
CAOD
Research Leader-Aircraft Performance, RL-Avionics
Simon Goss
Graeme Murray

DSTO Library

Library Fishermens Bend
Library Maribyrnong
Main Library DSTOS (2 copies)
Library, MOD, Pyrmont (Doc Data sheet)

Defence Central

OIC TRS, Defence Central Library
Officer in Charge, Document Exchange Centre (DEC), 1 copy
US Defence Technical Information Centre, 2 copies
UK Defence Research Information Centre, 2 copies
Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy
National Library of Australia, 1 copy
Defence Intelligence Organisation
Library, Defence Signals Directorate (Doc Data Sheet only)

Army

Director General Force Development (Land), (Doc Data)
ABCA Office, G-1-34, Russell Offices, Canberra (4 copies)

Navy

Director General Force Development (Sea), (Doc Data Sheet)

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Deakin University, Serials Section (M list)), Deakin University Library, Geelong, 3217
Senior Librarian, Hargrave Library, Monash University

OTHER ORGANISATIONS

NASA (Canberra)
AGPS
CSIRO DIT (Melbourne)
Australian Artificial Intelligence Institute, 3 copies
Prof. J. Hiller, Computer Science, University of NSW
Prof. T. Caelli, Computer Science, Curtin University
CRC Intelligent Division Systems
Prof J. Wallace, Swinburne University of Technology
Dr K. Bluff, Computer Science, Swinburne University of Technology

ABSTRACTING AND INFORMATION ORGANISATIONS

INSPEC: Acquisitions Section Institution of Electrical Engineers
Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
American Society for Metals
Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK
Library - Exchange Desk, National Institute of Standards and Technology, US

SPARES (10 copies)

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
				No	
2. TITLE Artificial Intelligence Applications in Aircraft Systems			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Simon Goss and Graeme Murray			5. CORPORATE AUTHOR Aeronautical and Maritime Research Laboratory PO Box 4331 Melbourne Vic 3001		
6a. DSTO NUMBER DSTO-RR-0071		6b. AR NUMBER AR-008-337		6c. TYPE OF REPORT Research Report	
7. DOCUMENT DATE February 1996					
8. FILE NUMBER M1/8/789	9. TASK NUMBER DST 92/091	10. TASK SPONSOR DSTO	11. NO. OF PAGES 59	12. NO. OF REFERENCES 23	
13. DOWNGRADING/DELIMITING INSTRUCTIONS			14. RELEASE AUTHORITY Chief, Air Operations Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT Approved for public release OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA ACT 2600					
16. DELIBERATE ANNOUNCEMENT No limitations					
17. CASUAL ANNOUNCEMENT Yes					
18. KEYWORD DESCRIPTORS Artificial Intelligence Air operations Cognitive modelling Aircombat modelling Trajectory optimisation Speech recognition Expert system					
19. ABSTRACT Air Operations Division at the DSTO Aeronautical and Maritime Research Laboratory is developing a capability in the use of Artificial Intelligence (AI), including knowledge based systems technology, in applications related to the operation and support of aircraft systems. A survey of the work program of Air Operations Division was undertaken to identify opportunities offered by advanced computing techniques for the solution of existing research problems. This document describes the findings of the survey. Some of the research opportunities identified have been pursued, and a brief description of progress is provided.					